

Title:

Intrinsic brain activity gradients dynamically coordinate functional connectivity states

Authors

Jesse A. Brown¹, Alex J. Lee¹, Lorenzo Pasquini¹, William W. Seeley¹

Affiliations

¹University of California, San Francisco, Memory and Aging Center,
Department of Neurology, San Francisco, CA, USA.

Corresponding Author

Jesse A. Brown, Ph.D.
Assistant Professor
Department of Neurology
University of California San Francisco
675 Nelson Rising Lane
San Francisco, CA 94158
jesse.brown@ucsf.edu

Abstract

Brain areas are organized into functionally connected networks though the mechanism underlying this widespread coordination remains unclear. Here we apply deep learning-based dimensionality reduction to task-free functional magnetic resonance images to discover the principal latent dimensions of human brain functional activity. We find that each dimension corresponds to a distinct and stable spatial activity gradient. Brain areas are distributed non-uniformly along each gradient, reflecting modular boundaries and hub properties. Gradients appear to dynamically steepen or flatten to produce task-specific activation patterns. Dynamical systems modelling reveals that gradients can interact via state-specific coupling parameters, allowing accurate forecasts and simulations of brain activity during different tasks. Together, these findings indicate that a small set of overlapping global activity gradients determine the repertoire of possible functional connectivity states.

Keywords

Gradients, Functional connectivity, Modularity, Deep learning, Dynamical systems

Introduction

Functional connectivity is defined as synchronous activity within two or more brain regions over time. Functional connectivity patterns as revealed by functional MRI (fMRI) are reliable within and across individuals (*Damoiseaux et al., 2006*) and have advanced our understanding of the brain's functional neuroanatomy in both health (*Power et al., 2011; Yeo et al., 2011*) and disease (*Seeley et al., 2009*). While the mechanism supporting functional connectivity remains unclear, the range of observed functional connectivity states appears diverse but bounded, both within individuals over short timescales (*Allen et al., 2014; Bassett et al., 2011; Calhoun et al., 2014; Lurie et al., 2020; Vidaurre et al., 2017*) and across individuals (*Gratton et al., 2018*). This constrained flexibility suggests that a “low-dimensional” set of factors may give rise to functional connectivity (*Glomb et al., 2019; MacDowell and Buschman, 2020; Saggar et al., 2018; Shine et al., 2019*). The neuroanatomical basis of these dimensions appear to be gradients – continuous, overlapping systems which reflect the predominant organization of brain regions into functionally connected networks (*Haak et al., 2018; Margulies et al., 2016*) and for which the structural underpinnings are becoming clearer (*Burt et al., 2018; Paquola et al., 2019; Wang, 2020*). A critical next step for validating the fundamental role of gradients in brain function is to demonstrate that individual gradients and their dynamic interactions can collectively explain key phenomena in functional connectomics. These include the presence of a brain-wide global signal (*Fox et al., 2009*), anti-correlation between large-scale networks (*Fox et al., 2005*), the presence of discrete functional modules and hubs regions (*Sporns and Betzel, 2016; van den Heuvel and Sporns, 2013*), correspondence with spatial patterns of gene expression (*Richiardi et al., 2015*), dynamic configuration into task-specific brain activity states (*Barch et al., 2013; Cole et al., 2014*), and stable functional connectivity patterns in individuals over time (*Finn et al., 2015*).

Low-dimensional brain activity latent space

We approached the study of gradient dynamics in task-free and task-engaged fMRI scans using deep neural networks, a powerful tool for dimensionality reduction (*Hinton and Salakhutdinov, 2006*) and the representation of spatial features (*Zeiler and Fergus, 2013*). We defined a latent space of functional brain activity where each dimension represents an independent spatial activity gradient and where the full space collectively represents the range of possible brain activity states. The latent space was learned from 119,500 task-free fMRI images, collected from a group of 100 healthy young control subjects in the Human Connectome Project (HCP; see **Methods**), using a 3D convolutional autoencoder and subsequent principal component

analysis (**Figure 1A**). By projecting each individual's scan into this low-dimensional space, we obtained latent trajectories of brain activity. To derive the spatial activity gradients associated with each latent dimension, we regressed each voxel's BOLD activity against that latent dimension timeseries, inferring a voxel's positive or negative weight (i.e. beta) on that dimension. A voxel's BOLD signal could be reconstructed by multiplying the voxel's weight on a gradient by the gradient's current slope and summing across all gradients (see **Figure 1C** and **Methods**). When reconstructing each brain region's timeseries, the first 3 dimensions explained 44.9% of BOLD variance across 273 cortical, subcortical, and cerebellar regions, while 9 dimensions explained 57.9% (**Figure S1** and **Supplementary Results**). This indicates that a low-dimensional latent space was able to explain a substantial proportion of the variation in BOLD activity, and therefore most subsequent analysis focused on these 9 dimensions.




Figure 1. Latent space derivation, gradient maps, and temporal interpretation.

A. The workflow for deriving individual latent trajectories and spatial activity gradient maps from task-free fMRI images. B. Activity gradient maps for the first 9 latent space dimensions. The sign of

each dimension is arbitrary. C. Illustration of relationship between latent space trajectories and regional BOLD activity. Across three successive timepoints, latent space positions on each dimension reflect the current slopes of the corresponding gradients. The resultant BOLD signal in each region depends on the region's weight on each gradient, shown here for the anterior cingulate cortex (yellow), premotor cortex (cyan), and precuneus (orange).

Each latent dimension was associated with a different but overlapping spatial gradient of brain activity. The spatial gradient weights for the first dimension were positive across 99.6% of the gray matter (**Figure 1B**, Gradient 1; the sign of each dimension is arbitrary), albeit with a topographically varied intensity that was highest in the primary sensory, visual, and auditory cortices. The gradient slope timeseries associated with this dimension had a near-perfect correlation with the global gray matter signal ($r=0.98$), a major influence on the estimated strength of functional connectivity (Fox et al., 2009). This was the only “unipolar” gradient, as no other gradient was more than 70% skewed towards positive- or negative-predominant weights. The spatial gradient for dimension 2 represented a sensory-to-cognitive axis (**Figure 1B**). The most positive weights were in areas of the default network (Greicius et al., 2003) and executive control network (Seeley et al., 2007) (see **Figure S4** for correspondence with 7 canonical networks (Yeo et al., 2011)), while the strongest negative weights were in somatomotor, visual, and ventral attention networks. The positioning of regions on the positive or negative side of a gradient signifies that these regions' gradient-specific activity will either be correlated or anti-correlated with each other (**Figure 1C**). This gradient had a strong correspondence with the putative primary functional connectivity gradient (Margulies et al., 2016) ($r=0.82$; see also **Supplementary Results**), reflecting a functional spectrum from perception and action to more abstract cognitive function. Dimension 3 resembled a task-positive (frontoparietal) to task-negative (default) gradient (**Figure 1B and Figure S4**). The subsequent dimensions included strongly lateralized activity gradients (Gradients 8 and 9), oppositions between specific sensory modalities like the visual and somatomotor networks (Gradients 4 and 5), and differential involvement of sub-components of larger super-systems like the default network (Gradients 2, 3, and 7) (Andrews-Hanna et al., 2010). Critically, gradients were highly reproducible in an independent validation dataset with the first 12 dimensions appearing in the nearly the same sequence (all spatial correlation coefficients $> r=0.5$; **Table S1 and Figure S4/Figure S5**). The consistency of the spatial gradient patterns across individuals

suggests that they may reflect a fundamental anatomical property of brain functional organization.

Basis for functional modularity and hubness

Building on the finding that gradients appear to be spatially consistent systems across individuals, we next characterized how the cumulative activity of each gradient at each timepoint can give rise to observed patterns of functional connectivity. As a demonstration of how to reach maximal activation of a region or co-variation of two regions based on gradient slopes, we conducted a whole-brain analysis using gradients 2 and 3 and then focused on two regions of interest, the anterior cingulate cortex (ACC) and middle frontal gyrus (MFG). For every region, we determined which latent space trajectory direction was optimal for most rapidly maximizing that region's BOLD activation relative to all other regions (**Figure 2A**). The ACC's activity was maximized when the slope of gradient 3 was more negative, reflected by the downward pointing arrows in Figure 2A (**top**). By contrast, the MFG's activity was maximized primarily by driving a more positive slope for gradient 2, reflected by the downward pointing arrows in Figure 2A (**bottom**). We next identified which latent trajectory angle maximized co-variation (i.e. functional connectivity) between the ACC and MFG. The optimal angle was from the top-left to bottom-right of this two-dimensional latent space, bisecting the trajectory angles that maximized activity in either the ACC or MFG. We confirmed that the individual subjects with maximal or minimal ACC-MFG functional connectivity during task-free fMRI had full scan-length latent trajectories that were most or least aligned with the optimal co-variation angle for this pair of regions (**Figure 2B**; see **Methods** and **Supplementary Results** for more detail). This demonstrates the value of rendering brain activity as a latent trajectory, providing a read out for levels of regional activity or functional connectivity for any timespan, from a single timepoint to a full scan. Importantly, individual subject scan-length latent trajectories (described by the gradient slope covariance matrix, see **Methods**) exhibited high reliability on consecutive days of scanning, consistent with the finding of identifiable individual functional connectivity "fingerprints" (Finn et al., 2015). 6 dimensions were required on average to correctly match a subject's day 1 and day 2 scans, while 29 dimensions were required to identify the correct fingerprint for all 100/100 subjects (**Supplementary Results** and **Figure S6**). Thus, a low dimensional latent space captures sufficient information to distinguish individuals, suggesting that reliable differences in gradient engagement may be an individual trait.




Figure 2. Functional modularity and hubness reflect non-uniform region gradient weights. **A.** Latent trajectory directions that maximize BOLD activation for the ACC, MFG, or the covariation between them. **B.** Latent trajectories during task-free scans for subjects with the most negative (top) or most positive (bottom) ACC-MFG correlation. **C.** Left: group-average functional connectivity matrix, based on region timeseries reconstructed only from latent dimensions 2 and 3. The 8 modules detected in this network are highlighted along the diagonal. Middle/Right: polar plots showing the angle of the optimal latent trajectory to maximize activation in each region, which canvases the full 360° of the two-dimensional latent space. Region colors are based on the modules they belong to, revealing that the bunching of regions with similar angles naturally reflects the modularity. Regions identified as the top connector hubs (middle) or provincial hubs (right) from the functional connectivity matrix have larger/smaller mean angles to their neighboring regions than do non-hubs.

The existence of spatially continuous activity gradients may appear to be at odds with the presence of discrete modular brain networks, a major principle of brain functional organization (*Sporns and Betzel, 2016*). We therefore attempted to reconcile the gradient and modular perspectives by testing the hypothesis that gradient non-uniformities would recapitulate modular boundaries. In the process of determining every region's optimal trajectory angle in a two-dimensional latent space for maximizing BOLD activation, we discovered

that the angles fully canvassed the 360° of the space (**Figure 2C**). Regions tended to cluster with their contralateral homologues and other regions in the same functional connectivity network, while regions that were diametrically opposed belonged to canonically anti-correlated networks. Based on this observation, we expected that functional connectivity modules derived from the functional connectome would correspond well with different angular ranges in latent space. We found that region module membership from the functional connectome corresponding exactly to the sequence of regions as grouped by optimal activation angle (**Figure 2C**). Consequently, provincial hubs and connector hubs (*van den Heuvel and Sporns, 2013*) were found to have characteristic orientations. Provincial hubs had significantly smaller angles to their most strongly connected neighbors than non-hubs (provincial hubs: mean=2.8°±2.3°; non-hubs: mean=4.0°±3.7°; F=80.45, p=5.72x10⁻¹⁹; **Figure S8**), which in turn had smaller angles to their neighbors than connector hubs (connector hubs: mean=4.8°±3.9°; F=9.94, p=0.001). This relationship held true when considering a higher dimensional latent space with a larger number of gradients (see **Supplemental Results**). Thus, the presence of modularity and hub regions appears to be consistent with the non-uniform distribution of regions along gradients. This finding helps unify the observations of both discrete, bounded networks (*Glasser et al., 2016*) and a graded functional-structural continuum.

Correspondence with spatial gene expression patterns

The spatial gradients described here and elsewhere have a striking congruence with the morphogen gradients that guide regional differentiation and connectivity during brain development (*O'Leary et al., 2007*). This motivated a systematic assessment by comparing each activity gradient's spatial similarity with genetic expression maps using the 15,655 genes across 196 cortical regions from the Allen Human Brain Atlas (*Fan et al., 2016; Hawrylycz et al., 2012*). Across the first 9 gradients, 4089 genes showed significant spatial correlations with at least one gradient in both the discovery and validation datasets, with correlation coefficients ranging between r=0.37–0.81 (surviving Bonferroni corrected threshold p < 3.55 x 10⁻⁷ in both the discovery and validation datasets; all correlations in **Supplementary Data 1**). The most striking correspondences were with gradient 1, for which 3572 genes were significantly correlated and which explained up to 64% of the variance in regional gradient weight. A gene ontology enrichment analysis based on the full set of significantly correlated genes found associations including "ion gated channel activity", "neuron projection", "synapse", and "anatomical structure development" (top-ranked terms in **Table S2**, all terms in **Supplementary Data 2**). The most significant positive relationships

between gene expression and gradient 1 were *SEMA7A* (discovery/validation mean $r=0.78$), *SCN1B* ($r=0.76$), *LAG3* ($r=0.76$), *ACAN* ($r=0.76$), *ASB13* ($r=0.76$), *SV2C* ($r=0.76$), *ANK1* ($r=0.76$), and *GPAT3* ($r=0.76$), while the most negative relationships were *KCNQ1* ($r=-0.81$), *ASCL2* ($r=-0.78$), *ANKRD6* ($r=-0.76$), *PRKCD* ($r=-0.75$), and *RSPH9* ($r=-0.75$). There was a significant spatial correlation between gradient 1 and the principal spatial component of gene expression ($r=0.72$), for which each of these individual genes was strongly loaded (all loading absolute Z scores > 2.8), revealing that the predominant sources of variability in BOLD activity and spatial gene expression are strongly linked. There were substantially fewer genes significantly correlated with the remaining gradients (**Supplementary Results** and **Figure S9**). Gradient 2 had 48 significantly correlated genes. Interestingly, gradient 3 had no significantly correlated genes (discovery minimum $p=2.51 \times 10^{-6}$, validation minimum $p=9.06 \times 10^{-11}$) despite explaining a substantial portion of BOLD variance, having a replicated spatial pattern in the validation dataset, and showing spatial correspondence to a previously described functional connectivity gradient (**Supplementary Results**). The lack of genetic correspondence for this gradient may be due to the stringent criterion for statistical significance or greater individual variability. Among the strongest of the other gene/gradient relationships were *CARTPT* on gradient 2 ($r=0.55$) and *CDH6*, *CDH13*, and *FABP6* on gradient 5 ($r=0.60/0.58/0.60$). Genes associated with activity gradients were recurrently linked to functional and structural processes likely to influence brain-wide activity patterns including anatomical morphogenesis, excitation/inhibition balance, and thalamocortical connectivity.

Gradient-based task brain activity patterns

We next evaluated the possibility that activity gradients detected during the task-free state are spatially stable, discrete systems that can dynamically adopt specific configurations to create task-specific brain activity states. We focused on four diverse cognitive tasks known to elicit distinct brain activity patterns – working memory (2-back vs. 0-back), motor movement (finger/toe/tongue movement vs. visual fixation), language comprehension (auditory story vs. math questions), and emotion processing (face emotion recognition vs. shape recognition). We projected HCP task fMRI data into the task-free state-defined latent space and derived gradient slope timeseries for each condition in each task. In each task, we found selective gradient slope differences between conditions ($\text{abs}(t) > 3.17$, $p < 0.005$) that combined to produce greater activation in areas consistent with previously reported task-activity patterns (**Figure 3**) (Barch et al., 2013). We verified the plausibility of the task-specific gradient-based activation patterns by measuring their spatial

correlation with conventional voxelwise GLM-based task activation maps. Gradient-based maps had strong correlation ($r > 0.5$) with the GLM-based maps when including between 3 to 10 gradients, rapidly increasing in correspondence before plateauing after ~20 gradients and reaching a maximum of $r=0.76-0.91$. The rapid increase and subsequent plateau of correlation suggests that the low-dimensional set of activity gradients enables diverse task-activation patterns.




Figure 3. Specific gradient slopes produce task-specific activity patterns.

A. Reconstructed BOLD activity differences between the active and baseline conditions in the working memory, motor, language, and emotion tasks, based on gradient slope differences of the first 9 latent dimensions. **B.** Gradient slope differences between the contrasting task conditions, where the line slope represents the across-subject average and the shaded band represents the 95% confidence interval. Black tick marks denote each region's weight on each gradient. **C.** Spatial correlation between the gradient-based task activation maps from panel A and the task activation maps derived with a conventional general linear model (GLM).

Predicting dynamic trajectories

Seeking evidence for the causal influence of activity gradients on brain activity dynamics, we tested their usefulness in forecasting and simulating brain activity states. Because latent space trajectories exhibited continuity and momentum, we elected to model the gradient slope timeseries as a dynamical system with differential equations representing the continuous influence of each activity gradient on one another (*Breakspear, 2017*). We used a data-driven strategy for parameter estimation (*Brunton et al., 2016*). For each fMRI volume, we measured the gradient slope (y), the 'velocity' of the slope (y'), and the 'acceleration' of the slope (y'') (**Figure 4A**). We then used a linear model for each of the first 9 gradients in the task-free state fMRI data to estimate instantaneous gradient slope acceleration as a function of all gradients' slopes and slope velocities. We found that a gradient's slope acceleration y'' primarily had a strong negative relationship with its own slope, (as is characteristic of an oscillating signal, ie $\sin(x)'' = -\sin(x)$), and secondarily with the slope velocity of the other 8 gradients (see **Supplementary Data 3**). Next, trajectories were modeled by using the parameter estimates from these linear models – the "coupling" parameters – to set up a system of second-order differential equations (**Figure 4A** and **Methods**). We measured the accuracy of these differential equations by using them to generate forecasts of task-free state gradient slope timeseries based on the initial slope and velocity of the 9 gradients at a given timepoint t . Critically, we used coupling parameters derived from the discovery dataset to make forecasts in the validation dataset. As a benchmark, we compared these forecasts to alternative forecasts generated by a first-order autoregressive model, a standard timeseries forecasting procedure which makes predictions based on the previous timepoint by leveraging the autocorrelation in the signal (*Liégeois et al., 2017*). The differential equation-based forecasts of the gradient slopes for gradients 1–9 at the subsequent fMRI timepoints explained significantly more variance than the autoregressive model (all $p < 0.001$): $t=2$, 99.8%/94.0% ($Z=385.0$); $t=3$, 96.9%/78.2% ($Z=244.5$); $t=4$, 87.3%/58.1% ($Z=162.5$); $t=5$, 69.2%/38.8% ($Z=105.2$); $t=6$, 46.1%/23.9% ($Z=60.7$); $t=7$, 24.7%/13.9% ($Z=26.9$); $t=8$, 10.3%/8.0% ($Z=5.5$). Thus, while forecasting brain activity at a long time horizon is not feasible, short-term predictions based on a model of dynamic interactions between activity gradients can outperform conventional timeseries forecasting methods.

A central tenet of cognitive neuroscience is that specific tasks will reliably evoke characteristic brain activity states. If specific modes of gradient coupling are required to induce task-specific

activation patterns, we expected that task-specific coupling parameters would generate more accurate forecasts of task fMRI activity timeseries than from other tasks or the task-free state. We tested this hypothesis in three ways: task-specific gradient timeseries reconstruction, simulation, and forecasting (**Figure 4A**). Reconstructions illustrated that the differential equations captured the within- and between-gradient factors that guide the moment-to-moment evolution of gradient changes for each task over a time horizon of ~20 timepoints (**Figure 4B, Figure S10, and Supplementary Results**). Next, we tested whether differential equations with state-specific coupling parameters could simulate timeseries that mimicked the brain-wide temporal dynamics in the actual task data as captured by task-specific functional connectivity patterns (see **Methods**). We found that for each task, the simulated timeseries functional connectivity matrix obtained with state-matched coupling parameters from the discovery dataset was always significantly more correlated to the true functional connectivity matrix from the same task condition than from any other task in the validation dataset (working memory, actual vs. simulated functional connectivity matrix $r=0.95$, $Z=12.09$; motor $r=0.94$, $Z=5.33$; language $r=0.93$, $Z=7.83$; emotion $r=0.95$, $Z=21.45$; all Z -associated $p < 0.001$; see **Figure 4C** for real and simulated functional connectivity matrices; **Figure 4D** for partial correlations among simulated matrices and actual matrices; see also **Table S3 and Supplementary Results**). To illustrate the influence of the coupling parameters on latent trajectories, we compared the shape of the simulated state-specific latent trajectories to the true latent trajectories in the working memory task (**Figure 4E**). The shapes and locations of the 2D trajectories were consistent across multiple dimensions, indicating that the simulations can replicate higher-order geometrical properties of the multi-dimensional data manifold. Finally, we performed timeseries forecasting to verify that state-matched coupling parameters could accurately extrapolate task-specific activity trajectories for unseen subjects. With parameters derived from the discovery dataset, the forecasts for all 4 tasks in the validation dataset were always significantly more accurate when based on the state-matched parameters than when using parameters from the baseline condition or from task-free state (working memory, for timepoints $t+1$ to $t+10$, all t -statistics versus baseline condition > 3 , all FDR-corrected $p < 0.01$; motor, timepoints $t+5$ to $t+10$; language, timepoints $t+1$ to $t+10$; emotion, timepoints $t+1$ to $t+10$; **Figure 4F and Supplementary Results**).




Figure 4. Gradient timeseries forecasts and simulations using coupled differential equations. A. Gradient forecasting schematic. For each gradient, a linear model estimated the gradient's acceleration as a function of all gradients' slopes and first derivatives. The

regression-derived coupling parameters were used to specify the system of coupled differential equations, which were then solved using initial conditions and coupling parameters that either come from in-sample data for illustrative purposes ("reconstructions"), or out of sample data for statistical assessment ("forecasts"). **B.** Differential equation-based reconstructions of gradient timeseries during the 2-back and 0-back conditions of the working memory task, based on trial-locked group-average data. The reconstructions accurately match the actual data up until $t=25$, after which they become unstable. FRNT: frontal, TEMP: temporal, PRT: parietal, INS: insula, CING: cingulate, OCCP: occipital, SUB: subcortical, CRB: cerebellar. **C.** Functional connectivity matrices from the active condition for each task, comparing actual data region pairwise correlations (upper triangle) with simulated data using state-specific coupling parameters (lower triangle). **D.** Partial correlations between each simulated functional connectivity matrix and the real/simulated functional connectivity matrices. **E.** Actual and simulated latent trajectories during working memory. Latent space difference histograms are shown for the first 9 latent dimensions, where actual working memory task trajectories on the upper triangle of the matrix can be compared to simulated trajectories on the lower triangle. Areas where trajectories dwell more during the 2-back than the 0-back condition are shown in red and vice versa in blue. **F.** Task fMRI forecast accuracy. In each task, coupling parameters from condition of interest (blue; *: significantly lower than red after FDR correction) yielded more accurate forecasts (root mean squared error \pm standard error of the mean) than task-free state parameters (green), or opposite task condition coupling parameters (red).

Discussion

Dimensionality reduction of brain activation images during task-free fMRI revealed a set of spatial brain activity gradients – overlapping, stable systems of influence on functional brain organization that dynamically adjust to enable a diverse repertoire of brain activation patterns. These gradients were consistently found in multiple groups of subjects, were reliable within individuals over time, present in both task-free and task-engaged states, and strongly correlated with brain gene expression patterns, comprehensively supporting their existence as stable anatomical systems. The spatial anatomy of the activity gradients and their collective interactions contribute to the field of human brain mapping by helping to unify a number of observed phenomena in fMRI: the global signal, gradients of functional connectivity, modularity, and the existence of hub regions. Dynamic reconfiguration of gradient temporal interactions appears

responsible for achieving task-specific brain activity states, offering a new perspective on the consistency of the brain's intrinsic architecture across rest and task (Smith et al., 2009).

The primary gradient in our study had a unipolar cortical activation pattern reflecting the global signal. A source of ongoing controversy in fMRI literature, the global signal appears to have a neuronal basis (Turchi et al., 2018) but has also been associated with respiration or head motion (Chang and Glover, 2009; Power et al., 2015). Here we found that this gradient had strongest involvement of early visual, somatomotor, and auditory areas, consistent with previous reports that the global signal has a heterogeneous spatial topography (Liu et al., 2018b). In line with this, this gradient had a strong spatial correlation with the principal spatial component of cortical gene expression, which has been linked to the cortical myelination pattern that demarcates the borders between sensory and association areas (Burt et al., 2018). The genes most positively associated with this gradient included *SEMA7A* and *SCN1B*, which have known roles in excitation/inhibition balance in somatomotor areas, thalamocortical projection, and seizure disorders including Dravet syndrome and genetic epilepsy with febrile seizures plus (Brackenburg et al., 2013; Carcea et al., 2014). Collectively, the primary gradient appears to represent a convergent system of functional, structural, and genetic variation in the brain which may modulate global neuronal excitation/inhibition balance.

The second gradient strongly resembled the principal macroscale gradient of functional connectivity (Margulies et al., 2016). This gradient defines a sensory-to-cognitive axis with the default network at one extreme and somatomotor/visual areas at the other. Balanced anticorrelation between networks is a central aspect of brain functional connectivity (Fox et al., 2005) and could plausibly be instantiated by placing brain areas at opposing ends of a single dynamic gradient. What circuit or systems-level mechanism might drive the ongoing fluctuations of these global, bipolar gradients? One compelling possibility is the reciprocal inhibitory connections in the thalamic reticular nucleus, which allow one thalamic nucleus to activate while inhibiting an opposing nucleus (Crabtree, 2018). This motif enables switching between attending to visual or auditory stimuli (Schmitt et al., 2017) and may enable thalamic coordination of widespread cortical functional connectivity (Halassa and Kastner, 2017; Hwang et al., 2017; Jones, 2001). Here, multiple gradients had correlations with genes including *SEMA7A* and *CDH6* that demarcate specific thalamic nuclei and are also expressed in cortical layers where thalamocortical axons terminate (Bibollet-Bahena et al., 2017; Bock and Goode, 2008; Nakagawa et al., 1999). Probing the link between microscale thalamic electrophysiology and macroscale BOLD activity gradient fluctuations may require an optogenetic fMRI approach (Liu et

al., 2015). A hypothetical second gradient switch may be in the amygdala, where two antagonistic, mutually inhibitory neuronal populations can drive either positive or negative valence behaviors (Kim et al., 2016). Intriguingly, one of these populations has been identified by CARTPT expression, the gene most strongly associated with our gradient 2.

One consequence of functional anticorrelation between systems is modularity, which the brain uses to perform segregated cognitive processing (Bertolero et al., 2015; Sporns and Betzel, 2016). Here we find that the presence of modular boundaries, provincial hubs, and connector hubs reflects the non-uniform distribution of regional weights on each functional activity gradient. This principle of non-uniformity or “clumpiness” has previously been observed to influence the distinctness of a region’s boundaries (Tian and Zalesky, 2018). Our observation that gradients had significant spatial correlations with gene expression maps implies that a non-uniform distribution of gene expression along gradients is also likely, with sections of a gradient expressing either clustered or transitioning gene expression profiles. Canonical functional connectivity networks have similarly been shown to correspond with spatial gene expression patterns (Bertolero et al., 2019; Hawrylycz et al., 2015; Richiardi et al., 2015). Morphogen gradients during brain development may provide the scaffold for the emergence of multiple distributed and overlapping activation gradients (O’Leary et al., 2007).

Our dynamical systems model shows that distinct brain activity states can be achieved by the dynamic interaction of stable anatomical gradients following state-specific coupling parameters. Specific modes of coupling may enable the brain to settle into transient stable states, where the amount of push and pull between gradients may be calibrated neuromodulation (Shine, 2019) or other mechanisms of gain control (Buzsáki, 2019). We observed a strong tendency for a transiently steep gradient slope to subsequently flatten out, suggesting that gradient engagement and maintenance may be energetically costly (Attwell and Laughlin, 2001). A tendency towards relaxation in parallel with mutual influence between gradients may lead to a “frustrated equilibrium” that perpetuates dynamic activity (Gollo and Breakspear, 2014). The shape of the latent brain activity manifold will be dictated by the coupling parameters, determining the coactivation patterns occurring during that state (Liu et al., 2018a) and the temporal sequence of activity flow (Cole et al., 2016). Future work can seek to understand how gradient coupling evolves on short timescales and is altered by neurological conditions or modified by feedback over extended timescales.

References

- Allen, E.A., Damaraju, E., Plis, S.M., Erhardt, E.B., Eichele, T., Calhoun, V.D., 2014. Tracking Whole-Brain Connectivity Dynamics in the Resting State. *Cereb. Cortex* 24, 663–676. <https://doi.org/10.1093/cercor/bhs352>
- Andrews-Hanna, J.R., Reidler, J.S., Sepulcre, J., Poulin, R., Buckner, R.L., 2010. Functional-anatomic fractionation of the brain's default network. *Neuron* 65, 550–562. <https://doi.org/10.1016/j.neuron.2010.02.005>
- Attwell, D., Laughlin, S.B., 2001. An energy budget for signaling in the grey matter of the brain. *J. Cereb. Blood Flow Metab* 21, 1133–1145. <https://doi.org/10.1097/00004647-200110000-00001>
- Barch, D.M., Burgess, G.C., Harms, M.P., Petersen, S.E., Schlaggar, B.L., Corbetta, M., Glasser, M.F., Curtiss, S., Dixit, S., Feldt, C., Nolan, D., Bryant, E., Hartley, T., Footer, O., Bjork, J.M., Poldrack, R., Smith, S., Johansen-Berg, H., Snyder, A.Z., Van Essen, D.C., 2013. Function in the human connectome: Task-fMRI and individual differences in behavior. *NeuroImage, Mapping the Connectome* 80, 169–189. <https://doi.org/10.1016/j.neuroimage.2013.05.033>
- Bassett, D.S., Wymbs, N.F., Porter, M.A., Mucha, P.J., Carlson, J.M., Grafton, S.T., 2011. Dynamic reconfiguration of human brain networks during learning. *Proc. Natl. Acad. Sci. U.S.A.* 108, 7641–7646. <https://doi.org/10.1073/pnas.1018985108>
- Bertolero, M.A., Blevins, A.S., Baum, G.L., Gur, R.C., Gur, R.E., Roalf, D.R., Satterthwaite, T.D., Bassett, D.S., 2019. The network architecture of the human brain is modularly encoded in the genome. *arXiv:1905.07606 [q-bio]*.
- Bertolero, M.A., Yeo, B.T.T., D'Esposito, M., 2015. The modular and integrative functional architecture of the human brain. *PNAS* 112, E6798–E6807. <https://doi.org/10.1073/pnas.1510619112>
- Bibollet-Bahena, O., Okafuji, T., Hokamp, K., Tear, G., Mitchell, K.J., 2017. A dual-strategy expression screen for candidate connectivity labels in the developing thalamus. *PLOS ONE* 12, e0177977. <https://doi.org/10.1371/journal.pone.0177977>
- Bock, G.R., Goode, J.A., 2008. *Cortical Development: Genes and Genetic Abnormalities*. John Wiley & Sons.
- Brackenbury, W.J., Yuan, Y., O'Malley, H.A., Parent, J.M., Isom, L.L., 2013. Abnormal neuronal patterning occurs during early postnatal brain development of *Scn1b*-null mice and precedes hyperexcitability. *Proc Natl Acad Sci U S A* 110, 1089–1094. <https://doi.org/10.1073/pnas.1208767110>
- Breakspear, M., 2017. Dynamic models of large-scale brain activity. *Nature Neuroscience* 20, 340–352. <https://doi.org/10.1038/nn.4497>
- Brunton, S.L., Proctor, J.L., Kutz, J.N., 2016. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *PNAS* 113, 3932–3937. <https://doi.org/10.1073/pnas.1517384113>
- Burt, J.B., Demirtaş, M., Eckner, W.J., Navejar, N.M., Ji, J.L., Martin, W.J., Bernacchia, A., Anticevic, A., Murray, J.D., 2018. Hierarchy of transcriptomic specialization across human cortex captured by structural neuroimaging topography. *Nature Neuroscience* 21, 1251–1259. <https://doi.org/10.1038/s41593-018-0195-0>

- Buzsáki, G., 2019. The Brain from Inside Out. Oxford University Press.
- Calhoun, V.D., Miller, R., Pearlson, G., Adalı, T., 2014. The chronnectome: time-varying connectivity networks as the next frontier in fMRI data discovery. *Neuron* 84, 262–274. <https://doi.org/10.1016/j.neuron.2014.10.015>
- Carcea, I., Patil, S.B., Robison, A.J., Mesias, R., Huntsman, M.M., Froemke, R.C., Buxbaum, J.D., Huntley, G.W., Benson, D.L., 2014. Maturation of cortical circuits requires Semaphorin 7A. *PNAS* 111, 13978–13983. <https://doi.org/10.1073/pnas.1408680111>
- Chang, C., Glover, G.H., 2009. Relationship between respiration, end-tidal CO₂, and BOLD signals in resting-state fMRI. *NeuroImage* 47, 1381–1393. <https://doi.org/10.1016/j.neuroimage.2009.04.048>
- Cole, M.W., Bassett, D.S., Power, J.D., Braver, T.S., Petersen, S.E., 2014. Intrinsic and Task-Evoked Network Architectures of the Human Brain. *Neuron* 83, 238–251. <https://doi.org/10.1016/j.neuron.2014.05.014>
- Cole, M.W., Ito, T., Bassett, D.S., Schultz, D.H., 2016. Activity flow over resting-state networks shapes cognitive task activations. *Nat. Neurosci.* 19, 1718–1726. <https://doi.org/10.1038/nn.4406>
- Crabtree, J.W., 2018. Functional Diversity of Thalamic Reticular Subnetworks. *Front. Syst. Neurosci.* 12. <https://doi.org/10.3389/fnsys.2018.00041>
- Damoiseaux, J.S., Rombouts, S. a. R.B., Barkhof, F., Scheltens, P., Stam, C.J., Smith, S.M., Beckmann, C.F., 2006. Consistent resting-state networks across healthy subjects. *PNAS* 103, 13848–13853. <https://doi.org/10.1073/pnas.0601417103>
- Fan, L., Li, H., Zhuo, J., Zhang, Y., Wang, J., Chen, L., Yang, Z., Chu, C., Xie, S., Laird, A.R., Fox, P.T., Eickhoff, S.B., Yu, C., Jiang, T., 2016. The Human Brainnetome Atlas: A New Brain Atlas Based on Connectional Architecture. *Cereb. Cortex* bhw157. <https://doi.org/10.1093/cercor/bhw157>
- Finn, E.S., Shen, X., Scheinost, D., Rosenberg, M.D., Huang, J., Chun, M.M., Papademetris, X., Constable, R.T., 2015. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nature Neuroscience* 18, 1664–1671. <https://doi.org/10.1038/nn.4135>
- Fox, M.D., Snyder, A.Z., Vincent, J.L., Corbetta, M., Van Essen, D.C., Raichle, M.E., 2005. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences of the United States of America* 102, 9673–9678. <https://doi.org/10.1073/pnas.0504136102>
- Fox, M.D., Zhang, D., Snyder, A.Z., Raichle, M.E., 2009. The Global Signal and Observed Anticorrelated Resting State Brain Networks. *J Neurophysiol* 101, 3270–3283. <https://doi.org/10.1152/jn.90777.2008>
- Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., Smith, S.M., Van Essen, D.C., 2016. A multi-modal parcellation of human cerebral cortex. *Nature* 536, 171–178. <https://doi.org/10.1038/nature18933>
- Glomb, K., Kringelbach, M.L., Deco, G., Hagmann, P., Pearson, J., Atasoy, S., 2019. Functional harmonics reveal multi-dimensional

- basis functions underlying cortical organization. *bioRxiv* 699678.
<https://doi.org/10.1101/699678>
- Gollo, L.L., Breakspear, M., 2014. The frustrated brain: from dynamics on motifs to communities and networks. *Philosophical Transactions of the Royal Society B: Biological Sciences* 369, 20130532.
<https://doi.org/10.1098/rstb.2013.0532>
- Gratton, C., Laumann, T.O., Nielsen, A.N., Greene, D.J., Gordon, E.M., Gilmore, A.W., Nelson, S.M., Coalson, R.S., Snyder, A.Z., Schlaggar, B.L., Dosenbach, N.U.F., Petersen, S.E., 2018. Functional Brain Networks Are Dominated by Stable Group and Individual Factors, Not Cognitive or Daily Variation. *Neuron* 98, 439–452.e5. <https://doi.org/10.1016/j.neuron.2018.03.035>
- Greicius, M.D., Krasnow, B., Reiss, A.L., Menon, V., 2003. Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences of the United States of America* 100, 253–258.
<https://doi.org/10.1073/pnas.0135058100>
- Haak, K.V., Marquand, A.F., Beckmann, C.F., 2018. Connectopic mapping with resting-state fMRI. *NeuroImage, Segmenting the Brain* 170, 83–94. <https://doi.org/10.1016/j.neuroimage.2017.06.075>
- Halassa, M.M., Kastner, S., 2017. Thalamic functions in distributed cognitive control. *Nature Neuroscience* 20, 1669.
<https://doi.org/10.1038/s41593-017-0020-1>
- Hawrylycz, M., Miller, J.A., Menon, V., Feng, D., Dolbeare, T., Guillozet-Bongaarts, A.L., Jegga, A.G., Aronow, B.J., Lee, C.-K., Bernard, A., Glasser, M.F., Dierker, D.L., Menche, J., Szafer, A., Collman, F., Grange, P., Berman, K.A., Mihalas, S., Yao, Z., Stewart, L., Barabási, A.-L., Schulkin, J., Phillips, J., Ng, L., Dang, C., Haynor, D.R., Jones, A., Van Essen, D.C., Koch, C., Lein, E., 2015. Canonical genetic signatures of the adult human brain. *Nature Neuroscience* 18, 1832–1844.
<https://doi.org/10.1038/nn.4171>
- Hawrylycz, M.J., Lein, E.S., Guillozet-Bongaarts, A.L., Shen, E.H., Ng, L., Miller, J.A., van de Lagemaat, L.N., Smith, K.A., Ebbert, A., Riley, Z.L., Abajian, C., Beckmann, C.F., Bernard, A., Bertagnolli, D., Boe, A.F., Cartagena, P.M., Chakravarty, M.M., Chapin, M., Chong, J., Dalley, R.A., David Daly, B., Dang, C., Datta, S., Dee, N., Dolbeare, T.A., Faber, V., Feng, D., Fowler, D.R., Goldy, J., Gregor, B.W., Haradon, Z., Haynor, D.R., Hohmann, J.G., Horvath, S., Howard, R.E., Jeromin, A., Jochim, J.M., Kinnunen, M., Lau, C., Lazarz, E.T., Lee, C., Lemon, T.A., Li, L., Li, Y., Morris, J.A., Overly, C.C., Parker, P.D., Parry, S.E., Reding, M., Royall, J.J., Schulkin, J., Sequeira, P.A., Slaughterbeck, C.R., Smith, S.C., Sodt, A.J., Sunkin, S.M., Swanson, B.E., Vawter, M.P., Williams, D., Wohnoutka, P., Zielke, H.R., Geschwind, D.H., Hof, P.R., Smith, S.M., Koch, C., Grant, S.G.N., Jones, A.R., 2012. An anatomically comprehensive atlas of the adult human brain transcriptome. *Nature* 489, 391–399.
<https://doi.org/10.1038/nature11405>
- Hinton, G.E., Salakhutdinov, R.R., 2006. Reducing the dimensionality of data with neural networks. *Science* 313, 504–507.
<https://doi.org/10.1126/science.1127647>

- Hwang, K., Bertolero, M.A., Liu, W.B., D'Esposito, M., 2017. The Human Thalamus Is an Integrative Hub for Functional Brain Networks. *J. Neurosci.* 37, 5594–5607. <https://doi.org/10.1523/JNEUROSCI.0067-17.2017>
- Jones, E.G., 2001. The thalamic matrix and thalamocortical synchrony. *Trends in Neurosciences* 24, 595–601. [https://doi.org/10.1016/S0166-2236\(00\)01922-6](https://doi.org/10.1016/S0166-2236(00)01922-6)
- Kim, J., Pignatelli, M., Xu, S., Itoharu, S., Tonegawa, S., 2016. Antagonistic negative and positive neurons of the basolateral amygdala. *Nat. Neurosci.* 19, 1636–1646. <https://doi.org/10.1038/nn.4414>
- Liégeois, R., Laumann, T.O., Snyder, A.Z., Zhou, J., Yeo, B.T.T., 2017. Interpreting temporal fluctuations in resting-state functional connectivity MRI. *NeuroImage* 163, 437–455. <https://doi.org/10.1016/j.neuroimage.2017.09.012>
- Liu, J., Lee, H.J., Weitz, A.J., Fang, Z., Lin, P., Choy, M., Fisher, R., Pinskiy, V., Tolpygo, A., Mitra, P., Schiff, N., Lee, J.H., 2015. Frequency-selective control of cortical and subcortical networks by central thalamus. *eLife* 4, e09215. <https://doi.org/10.7554/eLife.09215>
- Liu, X., Zhang, N., Chang, C., Duyn, J.H., 2018a. Co-activation patterns in resting-state fMRI signals. *NeuroImage, Brain Connectivity Dynamics* 180, 485–494. <https://doi.org/10.1016/j.neuroimage.2018.01.041>
- Liu, X., Zwart, J.A. de, Schölvinck, M.L., Chang, C., Ye, F.Q., Leopold, D.A., Duyn, J.H., 2018b. Subcortical evidence for a contribution of arousal to fMRI studies of brain activity. *Nature Communications* 9, 395. <https://doi.org/10.1038/s41467-017-02815-3>
- Lurie, D.J., Kessler, D., Bassett, D.S., Betzel, R.F., Breakspear, M., Kheilholz, S., Kucyi, A., Liégeois, R., Lindquist, M.A., McIntosh, A.R., Poldrack, R.A., Shine, J.M., Thompson, W.H., Bielczyk, N.Z., Douw, L., Kraft, D., Miller, R.L., Muthuraman, M., Pasquini, L., Razi, A., Vidaurre, D., Xie, H., Calhoun, V.D., 2020. Questions and controversies in the study of time-varying functional connectivity in resting fMRI. *Netw Neurosci* 4, 30–69. https://doi.org/10.1162/netn_a_00116
- MacDowell, C.J., Buschman, T.J., 2020. Low-Dimensional Spatio-Temporal Dynamics Underlie Cortex-Wide Neural Activity. *bioRxiv* 2020.01.05.895177. <https://doi.org/10.1101/2020.01.05.895177>
- Margulies, D.S., Ghosh, S.S., Goulas, A., Falkiewicz, M., Huntenburg, J.M., Langs, G., Bezgin, G., Eickhoff, S.B., Castellanos, F.X., Petrides, M., Jefferies, E., Smallwood, J., 2016. Situating the default-mode network along a principal gradient of macroscale cortical organization. *PNAS* 113, 12574–12579. <https://doi.org/10.1073/pnas.1608282113>
- Nakagawa, Y., Johnson, J.E., O'Leary, D.D.M., 1999. Graded and Areal Expression Patterns of Regulatory Genes and Cadherins in Embryonic Neocortex Independent of Thalamocortical Input. *J. Neurosci.* 19, 10877–10885. <https://doi.org/10.1523/JNEUROSCI.19-24-10877.1999>
- O'Leary, D.D.M., Chou, S.-J., Sahara, S., 2007. Area Patterning of the Mammalian Cortex. *Neuron* 56, 252–269. <https://doi.org/10.1016/j.neuron.2007.10.010>

- Paquola, C., Wael, R.V.D., Wagstyl, K., Bethlehem, R.A.I., Hong, S.-J., Seidlitz, J., Bullmore, E.T., Evans, A.C., Misic, B., Margulies, D.S., Smallwood, J., Bernhardt, B.C., 2019. Microstructural and functional gradients are increasingly dissociated in transmodal cortices. *PLOS Biology* 17, e3000284. <https://doi.org/10.1371/journal.pbio.3000284>
- Power, J.D., Cohen, A.L., Nelson, S.M., Wig, G.S., Barnes, K.A., Church, J.A., Vogel, A.C., Laumann, T.O., Miezin, F.M., Schlaggar, B.L., Petersen, S.E., 2011. Functional Network Organization of the Human Brain. *Neuron* 72, 665–678. <https://doi.org/10.1016/j.neuron.2011.09.006>
- Power, J.D., Schlaggar, B.L., Petersen, S.E., 2015. Recent progress and outstanding issues in motion correction in resting state fMRI. *NeuroImage* 105, 536–551. <https://doi.org/10.1016/j.neuroimage.2014.10.044>
- Richiardi, J., Altmann, A., Milazzo, A.-C., Chang, C., Chakravarty, M.M., Banaschewski, T., Barker, G.J., Bokde, A.L.W., Bromberg, U., Büchel, C., Conrod, P., Fauth-Bühler, M., Flor, H., Frouin, V., Gallinat, J., Garavan, H., Gowland, P., Heinz, A., Lemaître, H., Mann, K.F., Martinot, J.-L., Nees, F., Paus, T., Pausova, Z., Rietschel, M., Robbins, T.W., Smolka, M.N., Spanagel, R., Ströhle, A., Schumann, G., Hawrylycz, M., Poline, J.-B., Greicius, M.D., 2015. Correlated gene expression supports synchronous activity in brain networks. *Science* 348, 1241–1244. <https://doi.org/10.1126/science.1255905>
- Saggar, M., Sporns, O., Gonzalez-Castillo, J., Bandettini, P.A., Carlsson, G., Glover, G., Reiss, A.L., 2018. Towards a new approach to reveal dynamical organization of the brain using topological data analysis. *Nature Communications* 9, 1399. <https://doi.org/10.1038/s41467-018-03664-4>
- Schmitt, L.I., Wimmer, R.D., Nakajima, M., Happ, M., Mofakham, S., Halassa, M.M., 2017. Thalamic amplification of cortical connectivity sustains attentional control. *Nature* 545, 219–223. <https://doi.org/10.1038/nature22073>
- Seeley, W.W., Crawford, R.K., Zhou, J., Miller, B.L., Greicius, M.D., 2009. Neurodegenerative Diseases Target Large-Scale Human Brain Networks. *Neuron* 62, 42–52. <https://doi.org/10.1016/j.neuron.2009.03.024>
- Seeley, W.W., Menon, V., Schatzberg, A.F., Keller, J., Glover, G.H., Kenna, H., Reiss, A.L., Greicius, M.D., 2007. Dissociable intrinsic connectivity networks for salience processing and executive control. *J. Neurosci* 27, 2349–2356. <https://doi.org/10.1523/JNEUROSCI.5587-06.2007>
- Shine, J.M., 2019. Neuromodulatory Influences on Integration and Segregation in the Brain. *Trends in Cognitive Sciences* 0. <https://doi.org/10.1016/j.tics.2019.04.002>
- Shine, J.M., Breakspear, M., Bell, P.T., Martens, K.E., Shine, R., Koyejo, O., Sporns, O., Poldrack, R.A., 2019. Human cognition involves the dynamic integration of neural activity and neuromodulatory systems. *Nature Neuroscience* 1. <https://doi.org/10.1038/s41593-018-0312-0>
- Smith, S.M., Fox, P.T., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., Filippini, N., Watkins, K.E., Toro, R., Laird, A.R.,

- Beckmann, C.F., 2009. Correspondence of the brain's functional architecture during activation and rest. *Proceedings of the National Academy of Sciences* 106, 13040–13045. <https://doi.org/10.1073/pnas.0905267106>
- Sporns, O., Betzel, R.F., 2016. Modular Brain Networks. *Annu Rev Psychol* 67, 613–640. <https://doi.org/10.1146/annurev-psych-122414-033634>
- Tian, Y., Zalesky, A., 2018. Characterizing the functional connectivity diversity of the insula cortex: Subregions, diversity curves and behavior. *Neuroimage* 183, 716–733. <https://doi.org/10.1016/j.neuroimage.2018.08.055>
- Turchi, J., Chang, C., Ye, F.Q., Russ, B.E., Yu, D.K., Cortes, C.R., Monosov, I.E., Duyn, J.H., Leopold, D.A., 2018. The Basal Forebrain Regulates Global Resting-State fMRI Fluctuations. *Neuron* 97, 940–952.e4. <https://doi.org/10.1016/j.neuron.2018.01.032>
- van den Heuvel, M.P., Sporns, O., 2013. Network hubs in the human brain. *Trends in Cognitive Sciences, Special Issue: The Connectome* 17, 683–696. <https://doi.org/10.1016/j.tics.2013.09.012>
- Vidaurre, D., Smith, S.M., Woolrich, M.W., 2017. Brain network dynamics are hierarchically organized in time. *PNAS* 114, 12827–12832. <https://doi.org/10.1073/pnas.1705120114>
- Wang, X.-J., 2020. Macroscopic gradients of synaptic excitation and inhibition in the neocortex. *Nature Reviews Neuroscience* 21, 169–178. <https://doi.org/10.1038/s41583-020-0262-x>
- Yeo, B.T.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zöllei, L., Polimeni, J.R., Fischl, B., Liu, H., Buckner, R.L., 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology* 106, 1125–1165. <https://doi.org/10.1152/jn.00338.2011>
- Zeiler, M.D., Fergus, R., 2013. Visualizing and Understanding Convolutional Networks. *arXiv:1311.2901 [cs]*.

Acknowledgements

Funding

This work was supported by NIH grant K01-AG055698 to JAB and the Tau Consortium. Data was obtained from the Human Connectome Project (1U54MH091657, PIs Van Essen and Ugurbil).

Author Contributions: Jesse Brown: Conceptualization, Analysis, Writing. Alex Lee: Conceptualization, Analysis, Writing. Lorenzo Pasquini: Conceptualization, Analysis, Writing. William Seeley: Conceptualization, Writing.

Data and materials availability:

All code and data will be available upon publication at <https://github.com/jbrown81/gradients>