



Published in final edited form as:

Neuroimage. 2017 November 15; 162: 344–352. doi:10.1016/j.neuroimage.2017.08.042.

Instantaneous Brain Dynamics Mapped to a Continuous State Space

Jacob Billings¹, Alessio Medda², Sadia Shakil^{3,4}, Xiaohong Shen^{5,6}, Amrit Kashyap⁷, Shiyang Chen⁵, Anzar Abbas¹, Xiaodi Zhang⁵, Maysam Nezafati⁵, Wen-Ju Pan⁵, Gordon Berman⁸, and Shella Keilholz^{5,†}

¹Emory University, Graduate Division of Biological and Biomedical Sciences – Program in Neuroscience, Atlanta, USA

²Georgia Tech Research Institute, Division of Aerospace & Acoustics Technologies, Atlanta, USA

³Georgia Institute of Technology, School of Electrical and Computer Engineering, Atlanta, USA

⁴Institute of Space Technology, Department of Electrical Engineering, Islamabad, Pakistan

⁵Georgia Institute of Technology and Emory University, Department of Biomedical Engineering, Atlanta, USA

⁶Shandong University of Finance and Economics of China, School of Computer Science and Technology, Jinan, China

⁷Georgia Institute of Technology, School of Biological Engineering, Atlanta, USA

⁸Emory University, Department of Biology, Atlanta, USA

Abstract

Measures of whole-brain activity, from techniques such as functional Magnetic Resonance Imaging, provide a means to observe the brain's dynamical operations. However, interpretation of whole-brain dynamics has been stymied by the inherently high-dimensional structure of brain activity. The present research addresses this challenge through a series of scale transformations in the spectral, spatial, and relational domains. Instantaneous multispectral dynamics are first developed from input data via a wavelet filter bank. Voxel-level signals are then projected onto a representative set of spatially independent components. The correlation distance over the instantaneous wavelet-ICA state vectors is a graph that may be embedded onto a lower-dimensional space to assist the interpretation of state-space dynamics. Applying this procedure to a large sample of resting and task data (acquired through the Human Connectome Project), we segment the empirical state space into a continuum of stimulus-dependent brain states. We also demonstrate that resting brain activity includes brain states that are very similar to those adopted during some tasks, as well as brain states that are distinct from experimentally-defined tasks.

[†]Corresponding Author. Address 1760 Haygood Dr, Suite W-230; Atlanta, GA 30322-4600. shella.keilholz@bme.gatech.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Back-projection of segmented brain states onto the brain's surface reveals the patterns of brain activity that support each experimental state.

Keywords

fMRI; connectivity dynamics; functional connectivity; multiscale; dimensionality reduction

1. Introduction

The advent of functional Magnetic Resonance Imaging (fMRI) has launched the brain sciences into an exciting frontier by allowing the direct observation of systems-wide activity from healthy human brains (Rosen and Savoy, 2012). The richness of data this technology generates is the subject of cutting-edge research to interpret spontaneous signal fluctuations as indicators of preferential information exchange among the brain's intrinsic networks—i.e., its functional connectivity (FC) (Biswal et al., 1995; Hutchison et al., 2013). Brain FC networks were first defined over relatively long periods of time. Such *static* FC studies reveal that brain FC naturally develops a small-world topology, where densely connected local modules communicate with one another via richly interconnected hubs (Achard et al., 2006; Bullmore and Sporns, 2009). But the brain is not a static system. Rather, differential information exchange among neurons, circuits, and networks enable brains to deal flexibly with ever-changing environmental stimuli. The availability of rapid (< 1s), whole-brain imaging prompted researchers to look for shorter term *dynamics* of brain FC (Deco et al., 2011).

Early efforts to characterize brain dynamics observed that intra-network membership and internetwork communication possessed statistically significant differences when samples were drawn from short time windows during various epochs of an fMRI scan (Chang and Glover, 2010; Keilholz et al., 2013; Smith et al., 2012; Zalesky et al., 2014). While these short time window studies confirmed the expectation that the Blood-Oxygen Level Dependent (BOLD) fMRI signal may convey information about short-term brain-state dynamics, the large effect that *a priori* choices in window length had on study results lessened the method's analytic utility (Shakil et al., 2016). The effort to identify rapidly changing dynamics is also hampered by the drop-off in bold SNR at short window lengths.

To avoid the problems inherent in windowed analysis techniques, we present a method that provides a 2D map of the relative similarity of the brain's activity for all time points in the scan. The signal from each voxel first undergoes wavelet decomposition, making use of the BOLD signal's natural spectral scaling to characterize each time point as a summation of activations at multiple frequencies (Billings et al., 2015; Chang and Glover, 2010; Yaesoubi et al., 2015). This multispectral interpretation has been suggested to provide a parsimonious representation of the dynamic properties of complex systems like brains (Bullmore et al., 2004; Ciuciu et al., 2012; Mallat, 1989; Mandelbrot, 1983). To reduce the redundancy of spatial information and improve the SNR, voxel-wise signals are aggregated into a lower-dimensional spatial parcellation using Independent Component Analysis (ICA). In the

present study, we treat the collected vectors of multispectral activations from all of the ICA networks at each time point as samples of instantaneous brain states.

The dimensionality of the resulting data set is high (equal to the product of the number of functional networks and the number of spectral filters) and difficult to interpret. In order to explore the dynamics of brain activity, we apply t-distributed stochastic neighbor embedding (t-SNE) to represent the data from each time point in a two dimensional space (van der Maaten and Hinton, 2008), using correlation as a distance measure to ensure that similar states are grouped together. t-SNE is a state of the art data-driven dimensionality reduction algorithm that maintains local distance structure and has found wide application in the data-driven sciences to produce visualizations of *drosophila* behavior, machine learning hidden layers, static functional connectivity networks, and a host of other multidimensional structures (Berman et al., 2014; Mnih et al., 2015; Plis et al., 2014). In comparison to clustering based approaches that segment the time course into a number of predefined states, the map created by t-SNE produces a continuous distribution that can then be segmented empirically (using the watershed algorithm in this study). Information about the timing and the relative similarity of different states is preserved.

Towards the goal of detailing a map of brain-state dynamics, the present study analyzes the wide-ranging states 446 normal volunteers adopt as part of the Human Connectome Project (HCP)(Van Essen et al., 2012b). BOLD fMRI scans from 7 distinct tasks (EMOTION, GAMBLING, LANGUAGE, MOTOR, RELATIONAL, SOCIAL, and WORKING MEMORY (WM)), and from repeated resting conditions (REST1, and REST2) provide a basis to segment a t-SNE embedding of brain-state dynamics across experimentally defined events. We demonstrate the utility of the t-SNE mapping to characterize the human brain's coordination across time, space, and spectra during rest and in the negotiation of changing experimental stimuli.

2. Methods

Data Acquisition and Preprocessing

The data for this study was obtained from the HCP (Van Essen et al., 2012b). Whole-brain, BOLD-weighted, gradient-echo EPI data were acquired with a $TR = 0.720$ ms, and 2.0 mm isotropic voxels. Volunteers were scanned under 9 conditions, including: REST, EMOTION, GAMBLING, LANGUAGE, MOTOR, RELATIONAL, SOCIAL, and WORKING MEMORY (WM). The SOCIAL scan was examined in more detail during our analysis and is briefly described as follows: volunteers were presented 5 rounds of 20 s movies showing abstract objects making either random motions (*random*) or engaging in socially relevant movements (*mentalizing*). Each movie is followed by a 15 s fixation period where volunteers are asked to fix look at a '+' symbol. Each scan was performed twice.

A total of 446 volunteer datasets were included in the present study. Minimal data preprocessing was performed by HCP researchers. Steps included: spatial artifact and distortion removal, surface generation, anatomical registration, and alignment to grayordinate space. Voxel time series were normalized to zero mean and unit variance to fit the isotropic noise model expected by the ICA spatial filters. Each volunteer's fMRI data

was concatenated, across time, into a single matrix to minimize edge effects from spectral filtering. Scan order was randomized across volunteers.

Analysis

Previous studies have suggested that static functional connectivity networks segment into multiple frequency-specific architectures (Billings et al., 2015; Chang and Glover, 2010; Yaesoubi et al., 2015). Therefore, concatenated fMRI datasets were spectrally filtered into an octave of spectral bands, log-spaced over the low-frequency fluctuation range (0.1 to 0.01 Hz). Using the continuous wavelet transform schema, the filterbank was constructed from a low-order wavelet (Daubechies 4-tap wavelet) to provide optimal segmentation in the time domain with full coverage of the frequency domain (Daubechies, 1992). Brain images from each spectral band were multiplied by a 50 component group ICA spatial decomposition matrix. ICA filters were calculated as part of the HCP beta-release of group-ICA maps (Human Connectome Project, 2014). The number of components was chosen to just exceed the number needed for the eigenvalues of real and randomly shuffled data to be equal (data not shown). Time points were thus modeled as 400-dimensional states (8 spectral bands by 50 functional networks).

The state vectors for each time point were compared, pairwise, using the Pearson correlation distance. This choice highlights coordinated deviations from mean values. The correlation graph was then injected onto a 2-dimensional Euclidean surface using t-SNE. The t-SNE algorithm proceeds in two steps: first, the local neighborhood of each node is emphasized by normalizing inter-node distances via an adaptive Gaussian filter. Second, a 2-dimensional Euclidean version of the graph is constructed by minimizing the KL-divergence between the high-dimensional stochastic distribution and the low-dimensional stochastic distribution. A key innovation to t-SNE is to utilize a heavier tail in the low-dimensional probability distribution—a t-distribution rather than a Gaussian distribution. Doing so causes points that are only moderately far away in the high-dimensional space are pushed further apart in the low-dimensional representation. This feature allows for naturally affiliative clusters to emerge from an otherwise more compressed state space. Inherent similarities among sequentially sampled points will, none-the-less, cause these points to form their own distinct neighborhood. One way to encourage piecemeal-sequential points to arrive a group-level neighborhood is to embed points individually onto a group-level training embedding constructed from a sparse subsampling of each individual's scan data. We generated the training embedding in three steps. First, concatenated time series from each volunteer's full set of scans were t-SNE embedded into their own low-dimensional space. Second, ~2% (200 points) were selected from each volunteer's map to construct a group-level subsample. Third, the group-level subsample was t-SNE embedded to construct the training embedding. Subsequently, out-of-sample time points were injected onto the training embedding by satisfying the same symmetrized KL-divergence as used in to generate the training embedding. This subsampling procedure has the added benefit of reducing the computational load of embedding a large number of data points. (For additional details, please see van der Maaten and Hinton (2008), Berman and Shaevitz (2014), and the supplemental materials).

Quantitative Interpretation Methods

Density maps were constructed by convolving embedded point distributions by a 2-dimensional Gaussian filter. Density maps were compared using the structural similarity index (SSIM). SSIM is a robust measure of inter-image similarity as it is a product of three separate terms to account for salient image features, luminance, variance, and structure (Zhou et al., 2004). SSIM values range between 0 and 1, with 1 denoting identical images. Bootstrapped distributions of these comparisons provided sample statistics for ANOVA and multiple comparisons testing (see the section entitled, “Comparing Embeddings,” in the supplemental material).

To characterize the t-SNE embedding as disjoint brain states, the watershed transform of an (inverted) density map was used to establish boundaries around peaks (the catchment basins of the inverted image) (Meyer, 1994). The granularity of the watershed map is easily adjusted by changing the width of the Gaussian curve used to make the density maps. Given a discrete segmentation of brain states into watershed regions, a rudimentary count of the amount of time participants dwell in any given region is a measure of the stability of the regionally-bounded brain state. The map space is also amenable to group-level statistical testing. For example, we constructed a null model over the hypothesis that particular experimental conditions do not have a preference for any map region. Thereby, we investigated the preference that experimentally defined conditions have for the population of brain states embedded into each map region. Finally, brain state data were projected onto the midthickness Conte69 surface-based atlas to visualize anatomical characteristics (Van Essen et al., 2012a).

Additional methodological details, including the relevant equations, may be found in the supplemental information.

3. Results

To test the degree to which resting and tasked brains develop distinct dynamics, we segmented time points during the REST1 and REST2 scans from all task-scan time points. The results are displayed as density maps in part A of figure 1. The resting brain tends to adopt a range of states in the map’s periphery, while the task-active brain tends to develop brain states at the map’s interior. To represent the brain’s dynamic transitions across the embedded state space, part B displays point-to-point state changes as a velocity field. The results demonstrate that the resting brain’s most rapid transitions occur in regions densely populated during tasks. In the task-active segmentation, the highest velocities are found at the map’s center, between two interior regions densely populated during tasks. Regions of low velocity are distributed in patches throughout the task segmentation.

While the embedding is a continuous state space, the presence of multiple densely populated regions suggests an ensemble of discrete states that the brain adopts. Figure 2 takes on a discretized perspective by tracing boundaries around the resting-state segmentation’s dense regions. Formally, each state-space parcel is a catchment basin formed by taking the watershed transform of the density map’s inverse (Meyer, 1994). Part A of the figure codes regions in terms of the percentage of points in each parcel. Owing to its sheer size, a

sprawling domain in the map's interior contains the largest proportion of samples (4% to 5%, magenta boarder). The brain's propensity for adopting configurations within this parcel increases during task scans, when 6% to 7% of points form a similar density (see supplemental figure S1). The average brain state within this region sustains relatively slow (~0.019 Hz), low-amplitude, in-phase activations across most of the brain's static networks (part A, right) (see supplemental figure S2 for a description of each network). The region is often populated during the fixation periods of most tasks (see supplemental movies).

To gain a better understanding of the dynamic characteristics of resting-state parcels, part B of figure 2 displays the median time volunteers continuously dwelt in each parcel. Although maximum resting-state dwell times reached as high as 30 s (see supplemental figure S3), median dwell times did not exceed 2.5 s during rest. This finding is a confirmation that the resting brain often transitions between states. Dwell times tended to increase in duration during tasks (see supplemental figure S1). The mean brain state of one region having a long dwell time (magenta boarder) shows the brain to sustain activations in the same ICA networks as the state highlighted in part A. However, the second region's activations increase in magnitude. Further, they occur in two separate frequency bands, with either band's activations flipped to the opposite phase from the other.

The wide range of experimental states adopted during HCP scans provides a natural means to segment instantaneous brain-states. Figure 3 displays group-level map densities produced by such a segmentation. Bootstrap sampling provided a sample distribution to assess structural image similarity within-states, and also between-states (figure 4). Larger within-scan SSIM values indicate that the brain adopts a tighter range of states during the scan. Comparatively large between-scan SSIM values indicate that the two scans evoke similar varieties of brain states. Multiple comparisons statistics performed on these results determined that the repeated resting-state scans, alone, bear statistically similar continuous state distributions (see supplemental figure S4). Supplemental figure S5 displays the results of the same analysis when data are segmented against all block-design contrasts and task-related events.

The block design of HCP task scans—where stimuli are presented in rigidly timed sequences over several blocks—makes it possible to identify significant differences in the point-wise evolution of brain-states during the navigation of contrasting tasks, i.e., fine-scale brain dynamics. While the main text of the present manuscript uses the SOCIAL scan as an example, similar results are found for each set of block-design contrasts. (Movies illustrating the time-locked brain state distributions for all block-design contrasts may be found in the supplemental materials). Figure 5 outlines the point-wise state transitions during each SOCIAL task contrast's 35 s block. A bar chart of the mean SSIM at each time point surveys the focality of the progression of brain states evoked by either stimuli. Comparing the SSIM between stimuli provides a metric of state colocalization. The results demonstrate that, during the first 4 to 5 seconds of the stimulus, brain states in both conditions are incoherent. After ~7 s, both conditions achieve focal brain states, with the *mentalizing* condition being much more compactly delineated. After ~12 s, both conditions assume similar brain states. The onset of the fixation block causes brain states to, again, disperse across the embedding. In the latter half of the fixation block, participants who witnessed a *mentalizing* condition

tend to linger in a very tightly localized brain state. This suggests that a more characteristic variety of rumination occurs in response to the *mentalizing* stimulus. Density maps from time points having focal brain states are shown in the insets of figure 5 and close up in figure 6.

To check the uniqueness of each contrast's associated brain states, figure 6 magnifies the most densely populated regions in each contrast's embedding at those time points presented in figure 5. After ~8 s, volunteers' brains are observed to adopt adjacent and disjoint states. A short time later (~15 s), participants may adopt similar states albeit with the *mentalizing* stimulus inducing brains to adopt a more focal subset of the *random* stimulus' state space. A similar observation obtains during the fixation block (~32 s) with the *mentalizing* stimuli evoking a focal subset relative to the *random* stimuli's state space.

The extremely focused distribution of brain states during particular moments of block-design tasks motivated a closer investigation of the statistical distribution of embedded points at particular moments in time. To conduct this analysis, we generated a watershed segmentation of the embedded space after convolving all points on the embedded map with a very narrow Gaussian filter (figure 7, part A). We then labeled each embedded point in terms of the experimental condition under which the brain-state was generated, as well as in terms of the time that state was generated relative to the start of each experimental block. Next, we randomly permuted the point labels 100 times to generate a null distribution of the embedded point locations for each condition, at each time point. Finally, we calculated the z-statistic of the probability that the number of embedded points in each watershed region was greater in the real data than in the permuted data. The significance threshold was initially set to a p-value of 5%. With Bonferroni correction for multiple comparisons across ~5000 watershed regions and ~1000 individual time points, the significance threshold was set to a p-values less than $1e-8$. Part B of figure 7 color-codes fine-grained watershed regions in terms of the most probable state associated with that region. For simplicity, all time points from a given condition share the same color coding. Part C of the figure addresses the inference from figures 5 and 6 that the contrasting conditions in the SOCIAL task result in highly stereotypical brain states at especially $t \cong 8s$ after the start of the block. As inferred from figure 6, the *social* stimulus induces highly focal brain states in the map space below and to the right of grid location e3 (green borders). While some brain states generated by the *random* stimulus overlap this region (yellow border), these brain states mostly lie above grid line e or to the left of grid line 3 (red borders).

We can gain insights into how the brain responds to the contrasting stimuli by projecting a local averaging of the state space (magenta points) onto a model brain surface. Figure 8 displays the contrast between the *mentalizing* and *random* stimuli at 7.92 s (TR = 11). Whereas higher frequency (≥ 0.037 Hz) activations are similar, the lower frequency (≤ 0.019 Hz) brain states bear marked differences. At infra-slow frequencies (0.01 Hz) the *mentalizing* stimuli induces in-phase oscillations between the visual, parietal, sensorimotor, and lateral prefrontal cortices. This contrasts with brains experiencing the *random* stimulus for which the visual and left parietal cortex are out-of-phase relative to the anterior prefrontal, left orbitofrontal, and left parietal networks. Slow (0.019 Hz) activations are similar between block contrasts save for, 1) the inclusion of the medial prefrontal cortex

within the positive-phase network during the *mentalizing* condition, and 2) stronger negative-phase activation in the anterior prefrontal cortex among volunteers receiving the *random* stimulus. Task-based activation studies of the same task-contrasts identified similar areas of contrasting brain activity, including the medial prefrontal cortex, lateral parietal cortices, and the visual cortex (Barch et al., 2013; Castelli et al., 2000).

4. Discussion

The present analytical framework, where BOLD dynamics are interpreted as multiscalar, instantaneous events overcomes many of the challenges faced in the study of brain dynamics. Unlike methods based on sliding window correlation, it avoids the challenges involved in choosing a window length (Hindriks et al., 2016; Shakil et al., 2016). Additionally, whereas previous studies using clustering tended to delineate brain states into a fixed number, k , of categories (Calhoun et al., 2014), manifold embedding optimizes a low-dimensional representation of the high-dimensional data, allowing a continuous distribution. Nevertheless, it remains possible to identify discrete state categories via subsequent analyses (e.g., through a watershed transformation of the continuous state space).

The statistically significant differences between the state-space distributions of each task provides assurance that the embedded state space effectively differentiates between activation patterns related to different tasks (figures 3 and 4). On the other hand, the statistically insignificant differences between the state-space distributions of the repeated resting-state scans provide assurance that the embedded state space does not overspecify differences between brain states. The identification of fine-grained differences between the brain states of task contrasts confirms our hypothesis that common stimuli result in short distances between embedding points. Our qualitative and quantitative analysis shows that the random and mentalizing portions of the social task inhabit adjacent but disjoint map regions (figures 5 through 8).

The segmentation of the embedded space allows the identification of networks whose coactivations (or lack there-of), at particular frequencies and phases, are predominant in any given state (figure 2). Our analysis of the SOCIAL scan demonstrates this point when brain regions involved in attention (lateral prefrontal) mental representation (parietal cortex) and somatic representation (somatomotor cortex) are slowly driven, in-phase, with activations of the visual system.

One of the advantages of the 2D representation obtained from t-SNE is the ease of exploratory data analysis. From the distribution of the data, one can hypothesize about the similarities of different tasks, identify common trajectories, identify states and substates, etc. The resulting hypotheses can be addressed through further statistical analysis as demonstrated for the social task, or they might motivate more specific experiments designed to address the questions in other ways. Regardless of the following analysis, t-SNE provides a powerful tool for characterizing functional neuroimaging data.

Insight into resting state fMRI

Given the success of t-SNE at differentiating between tasks, the present study's embedded state space may offer new insights into lingering questions on the character of the resting state (Lowe, 2012). The wide spatial extent, absence of low velocity regions, and overall short dwell times, converge on the finding that the resting-state is not a singular condition. As the only difference between rest and task is the absence of an explicit stimuli, the preference resting brains display for peripheral map regions marks the resting-state as mostly distinct from each of the 7 task states. One notable exception is the interior map region roughly bounded by grid vertices f5, f7, e5 and e7. Details of the region's contribution from each task contrast (supplemental figure 5) the LANGUAGE scan's *story* condition, the EMOTION scan's *neutral* stimulus, WM *0 back* challenges, and the collection of time points when no stimulus information was explicitly provided (labeled rest). The region may therefore relate to times when volunteers are externally oriented albeit with low cognitive demands. Indeed, volunteer brains often populate this region during the fixation blocks of most tasks (see supplemental videos of especially the GAMBLING and RELATIONAL scans). Another notable exception is the projection of MOTOR and SOCIAL brain states to locations further out in the map's periphery than the significant REST regions in figure 7 part B. Despite the general lack of overlap between the REST and TASK conditions, the REST condition itself exhibits several dense concentrations of time points similar to those observed during the tasks, suggesting that these network configurations constitute common brain states during the resting condition.

Limitations

The generation of the embedding's features comes directly from a combination of the input data and the analytical model. As always, the BOLD signal's lack of direct sensitivity to neural activity limits our ability to infer the underlying neurophysiology from the functional imaging data. Noise in the BOLD signal (from residual physiological noise, motion, or the scanner) will affect the embedding of the data.

Regarding the analysis itself, the symmetric distribution of contrasting brain-states across the map may owe itself to the use of signed wavelet coefficients. Like all spectral decompositions, the wavelet transform inherently generates phase information as complex coefficients. The present study projects complex coefficients onto the set of real numbers, thus limiting the analysis to account for two phases, separated by 180°. Other studies have found good segmentation when comparing BOLD signals across additional phases (Chang and Glover, 2010; Yaesoubi et al., 2015). The utility of incorporating phase information supports the notion that regional activations bear some degree of phase-coupling (Thompson GJ, 2014; Tort et al., 2010). Future studies may appropriate this natural feature of brain activity by characterizing the data in alternative metric spaces that better utilize complex-valued data in the high-dimensional space, and better distribute their states into a low-dimensional space.

Furthermore, it should be noted that while using as a wavelet kernel for the spectral transform is expected to result in improved time frequency localization relative to a short-time Fourier transform, it is possible to tune the kernel function to extract additional

information about BOLD dynamics. In addition to requiring that the transform kernel decay smoothly to 0 after a short time, we might also require that the kernel somehow better represent the BOLD signal's underlying impulse waveform. Such kernels may be developed by directly "lifting" the temporal shape from the data itself (Sweldens, 1998). Such a procedure would provide better localization of the BOLD signal's energy into fewer wavelet coefficients, and thus improve the ability to differentiate between states.

Future directions

This study demonstrates an analytical technique to observe BOLD dynamics that performs well in segmenting contrasting activities. This approach provides a simple way to summarize patterns of brain state dynamics across various tasks as well as when participants are at rest. The ready capacity to chart brain-state dynamics against experimental stimuli raises the interest for demarcating the functional space of other varieties of conditions. Indeed, study methods are readily amenable to describing brain state dynamics of differing populations and animal models including patient populations. One potential application for the t-SNE embedding is to determine whether it can identify specific states that are present in patients but not in healthy controls. Another area of interest utilizes the preserved timing information to determine common trajectories of brain activity across states during task and rest. The t-SNE embedding facilitates exploratory analysis but can also be used to identify significant differences between tasks or populations using additional analysis. While permutation was applied for most of the statistical analyses shown in this manuscript, more sophisticated approaches should also be pursued.

The subsampling procedure reduces the computational complexity of fitting an out of sample point from $O(n^2)$ operations to $O(n)$. This feature enables future research to chart increasingly detailed and comprehensive maps of the brain's dynamical state space from an ever-increasing pool of shared data. One future area of investigation should examine whether data acquired on different scanners or with different parameters (TR, for example) can be added to the existing embedding or should be handled separately.

5. Conclusion

The BOLD signal's multispectral components, developed naturally among the brain's many networks, constitute a robust descriptor of instantaneous brain states. High-dimensional datasets containing point-to-point brain state dynamics are made interpretable by embedding the graph onto a 2-dimensional sheet. Our analysis of a dynamical brain-state embedding from a large population ($N=446$) concludes that the resting brain actively pursues a range of distinctive states from those adopted during explicit tasks. The realization of both resting and task-active states involves large-scale, and often phase-locked coordination's among multiple brain regions at particular frequencies. The tools described in this manuscript should enable further analysis of the brain's dynamic trajectories during both tasks and rest.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors gratefully acknowledge Dr. Dieter Jaeger and the Emory Center for Mind, Brain, and Culture whose *Dimensionality Reduction Workshop* inspired this work. Special thanks go to Dr. Ying Guo who suggested the permutation test for statistical significance. These efforts were directly funded by: NIH 5-R01NS078095-02, and by Professional Development Supports Funds provided by the Laney Graduate School at Emory University. Data were provided by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

References

- Achard S, Salvador R, Whitcher B, Suckling J, Bullmore E. A Resilient, Low-Frequency, Small-World Human Brain Functional Network with Highly Connected Association Cortical Hubs. *The Journal of Neuroscience*. 2006; 26:63–72. [PubMed: 16399673]
- Barch DM, Burgess GC, Harms MP, Petersen SE, Schlaggar BL, Corbetta M, Glasser MF, Curtiss S, Dixit S, Feldt C, Nolan D, Bryant E, Hartley T, Footer O, Bjork JM, Poldrack R, Smith S, Johansen-Berg H, Snyder AZ, Van Essen DC. Function in the Human Connectome: Task-fMRI and Individual Differences in Behavior. *Neuroimage*. 2013; 80:169–189. [PubMed: 23684877]
- Berman GJ, Choi DM, Bialek W, Shaevitz JW. Mapping the stereotyped behaviour of freely moving fruit flies. *Journal of The Royal Society Interface*. 2014;11.
- Billings, JC., Medda, A., Keilholz, SD. Multiscale FC analysis refines functional connectivity networks in individual brains; 2015 IEEE Global Conference on Signal and Information Processing (GlobalSIP); 2015. p. 557-561.
- Biswal B, Yetkin FZ, Haughton VM, Hyde JS. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn Reson Med*. 1995; 34:537–541. [PubMed: 8524021]
- Bullmore E, Fadili J, Maxim V, Sendur L, Whitcher B, Suckling J, Brammer M, Breakspear M. Wavelets and functional magnetic resonance imaging of the human brain. *Neuroimage*. 2004; 23(Suppl 1):S234–249. [PubMed: 15501094]
- Bullmore E, Sporns O. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci*. 2009; 10:186–198. [PubMed: 19190637]
- Calhoun, Vince D., Miller, R., Pearlson, G., Adalı, T. The Chronnectome: Time-Varying Connectivity Networks as the Next Frontier in fMRI Data Discovery. *Neuron*. 2014; 84:262–274. [PubMed: 25374354]
- Castelli F, Happé F, Frith U, Frith C. Movement and Mind: A Functional Imaging Study of Perception and Interpretation of Complex Intentional Movement Patterns. *Neuroimage*. 2000; 12:314–325. [PubMed: 10944414]
- Chang C, Glover GH. Time-frequency dynamics of resting-state brain connectivity measured with fMRI. *Neuroimage*. 2010; 50:81–98. [PubMed: 20006716]
- Ciuciu P, Varoquaux G, Abry P, Sadaghiani S, Kleinschmidt A. Scale-Free and Multifractal Time Dynamics of fMRI Signals during Rest and Task. *Front Physiol*. 2012; 3:186. [PubMed: 22715328]
- Daubechies, I. Society for Industrial and Applied Mathematics. Philadelphia, PA: 1992. Ten lectures on wavelets.
- Deco G, Jirsa VK, McIntosh AR. Emerging concepts for the dynamical organization of resting-state activity in the brain. *Nat Rev Neurosci*. 2011; 12:43–56. [PubMed: 21170073]
- He BJ. Scale-free brain activity: past, present, and future. *Trends Cogn Sci*. 2014; 18:480–487. [PubMed: 24788139]
- Hindriks R, Adhikari MH, Murayama Y, Ganzetti M, Mantini D, Logothetis NK, Deco G. Can sliding-window correlations reveal dynamic functional connectivity in resting-state fMRI? *Neuroimage*. 2016; 127:242–256. [PubMed: 26631813]
- Human Connectome Project. HCP beta-release of group-ICAmaps, nodetimeseries and network matrices. 2014. <http://humanconnectome.org>
- Hutchison RM, Womelsdorf T, Allen EA, Bandettini PA, Calhoun VD, Corbetta M, Della Penna S, Duyn JH, Glover GH, Gonzalez-Castillo J, Handwerker DA, Keilholz S, Kiviniemi V, Leopold

- DA, de Pasquale F, Sporns O, Walter M, Chang C. Dynamic functional connectivity: promise, issues, and interpretations. *Neuroimage*. 2013; 80:360–378. [PubMed: 23707587]
- Keilholz SD, Magnuson ME, Pan WJ, Willis M, Thompson GJ. Dynamic properties of functional connectivity in the rodent. *Brain Connect*. 2013; 3:31–40. [PubMed: 23106103]
- Liu X, Ward BD, Binder JR, Li S-J, Hudetz AG. Scale-Free Functional Connectivity of the Brain Is Maintained in Anesthetized Healthy Participants but Not in Patients with Unresponsive Wakefulness Syndrome. *PLoS ONE*. 2014; 9:e92182. [PubMed: 24647227]
- Lowe MJ. The emergence of doing “nothing” as a viable paradigm design. *Neuroimage*. 2012; 62:1146–1151. [PubMed: 22245648]
- Mallat SG. A theory for multiresolution signal decomposition: the wavelet representation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*. 1989; 11:674–693.
- Mandelbrot, BB. *The fractal geometry of nature*. Macmillan: 1983.
- Meyer F. Topographic distance and watershed lines. *Signal Processing*. 1994; 38:113–125.
- Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G, Petersen S, Beattie C, Sadik A, Antonoglou I, King H, Kumaran D, Wierstra D, Legg S, Hassabis D. Human-level control through deep reinforcement learning. *Nature*. 2015; 518:529–533. [PubMed: 25719670]
- Plis SM, Hjelm D, Salakhutdinov R, Allen EA, Bockholt HJ, Long JD, Johnson HJ, Paulsen J, Turner JA, Calhoun VD. Deep learning for neuroimaging: a validation study. *Front Neurosci*. 2014;8. [PubMed: 24550770]
- Rosen BR, Savoy RL. fMRI at 20: Has it changed the world? *Neuroimage*. 2012; 62:1316–1324. [PubMed: 22433659]
- Shakil S, Lee CH, Keilholz SD. Evaluation of sliding window correlation performance for characterizing dynamic functional connectivity and brain states. *Neuroimage*. 2016; 133:111–128. [PubMed: 26952197]
- Smith SM, Miller KL, Moeller S, Xu J, Auerbach EJ, Woolrich MW, Beckmann CF, Jenkinson M, Andersson J, Glasser MF, Van Essen DC, Feinberg DA, Yacoub ES, Ugurbil K. Temporally-independent functional modes of spontaneous brain activity. *Proceedings of the National Academy of Sciences*. 2012; 109:3131–3136.
- Sweldens W. The Lifting Scheme: A Construction of Second Generation Wavelets. *SIAM Journal on Mathematical Analysis*. 1998; 29:511–546.
- Thompson GJPW, Billings JC, Grooms JK, Shakil S, Jaeger D, Keilholz SD. Phase-amplitude coupling and infraslow (<1Hz) frequencies in the rat brain: relationship to resting state fMRI. *Front. Integr. Neurosci*. 2014;8. [PubMed: 24605091]
- Tort ABL, Komorowski R, Eichenbaum H, Kopell N. Measuring Phase-Amplitude Coupling Between Neuronal Oscillations of Different Frequencies. *Journal of Neurophysiology*. 2010; 104:1195–1210. [PubMed: 20463205]
- van der Maaten L, Hinton G. Visualizing data using t-SNE. *Journal of Machine Learning Research*. 2008; 9:85.
- van der Maaten L, Postma E, Van den Herik J. Dimensionality reduction: a comparative. *J Mach Learn Res*. 2009; 10:66–71.
- Van Essen DC, Glasser MF, Dierker DL, Harwell J, Coalson T. Parcellations and hemispheric asymmetries of human cerebral cortex analyzed on surface-based atlases. *Cereb Cortex*. 2012a; 22:2241–2262. [PubMed: 22047963]
- Van Essen DC, Ugurbil K, Auerbach E, Barch D, Behrens TE, Bucholz R, Chang A, Chen L, Corbetta M, Curtiss SW, Della Penna S, Feinberg D, Glasser MF, Harel N, Heath AC, Larson-Prior L, Marcus D, Michalareas G, Moeller S, Oostenveld R, Petersen SE, Prior F, Schlaggar BL, Smith SM, Snyder AZ, Xu J, Yacoub E. The Human Connectome Project: a data acquisition perspective. *Neuroimage*. 2012b; 62:2222–2231. [PubMed: 22366334]
- Wornell GW. Wavelet-based representations for the 1/f family of fractal processes. *Proceedings of the IEEE*. 1993; 81:1428–1450.
- Yaesoubi M, Allen EA, Miller RL, Calhoun VD. Dynamic coherence analysis of resting fMRI data to jointly capture state-based phase, frequency, and time-domain information. *Neuroimage*. 2015; 120:133–142. [PubMed: 26162552]

- Zalesky A, Fornito A, Cocchi L, Gollo LL, Breakspear M. Time-resolved resting-state brain networks. *Proceedings of the National Academy of Sciences*. 2014; 111:10341–10346.
- Zhou W, Bovik AC, Sheikh HR, Simoncelli EP. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*. 2004; 13:600–612. [PubMed: 15376593]

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Highlights

- We demonstrate the construction and interrogation of a continuous, two-dimensional map of fMRI dynamics.
- Map points represent an individual's multispectral, and multispectral BOLD state centered at a single point in time.
- Task-based scans occupy focal state-spaces, reinforcing the utility of study methods to capture salient BOLD dynamics evoked by experimental stimuli.
- Resting-state scans occupy a broad state-space, reinforcing the view that the resting mind is highly active.

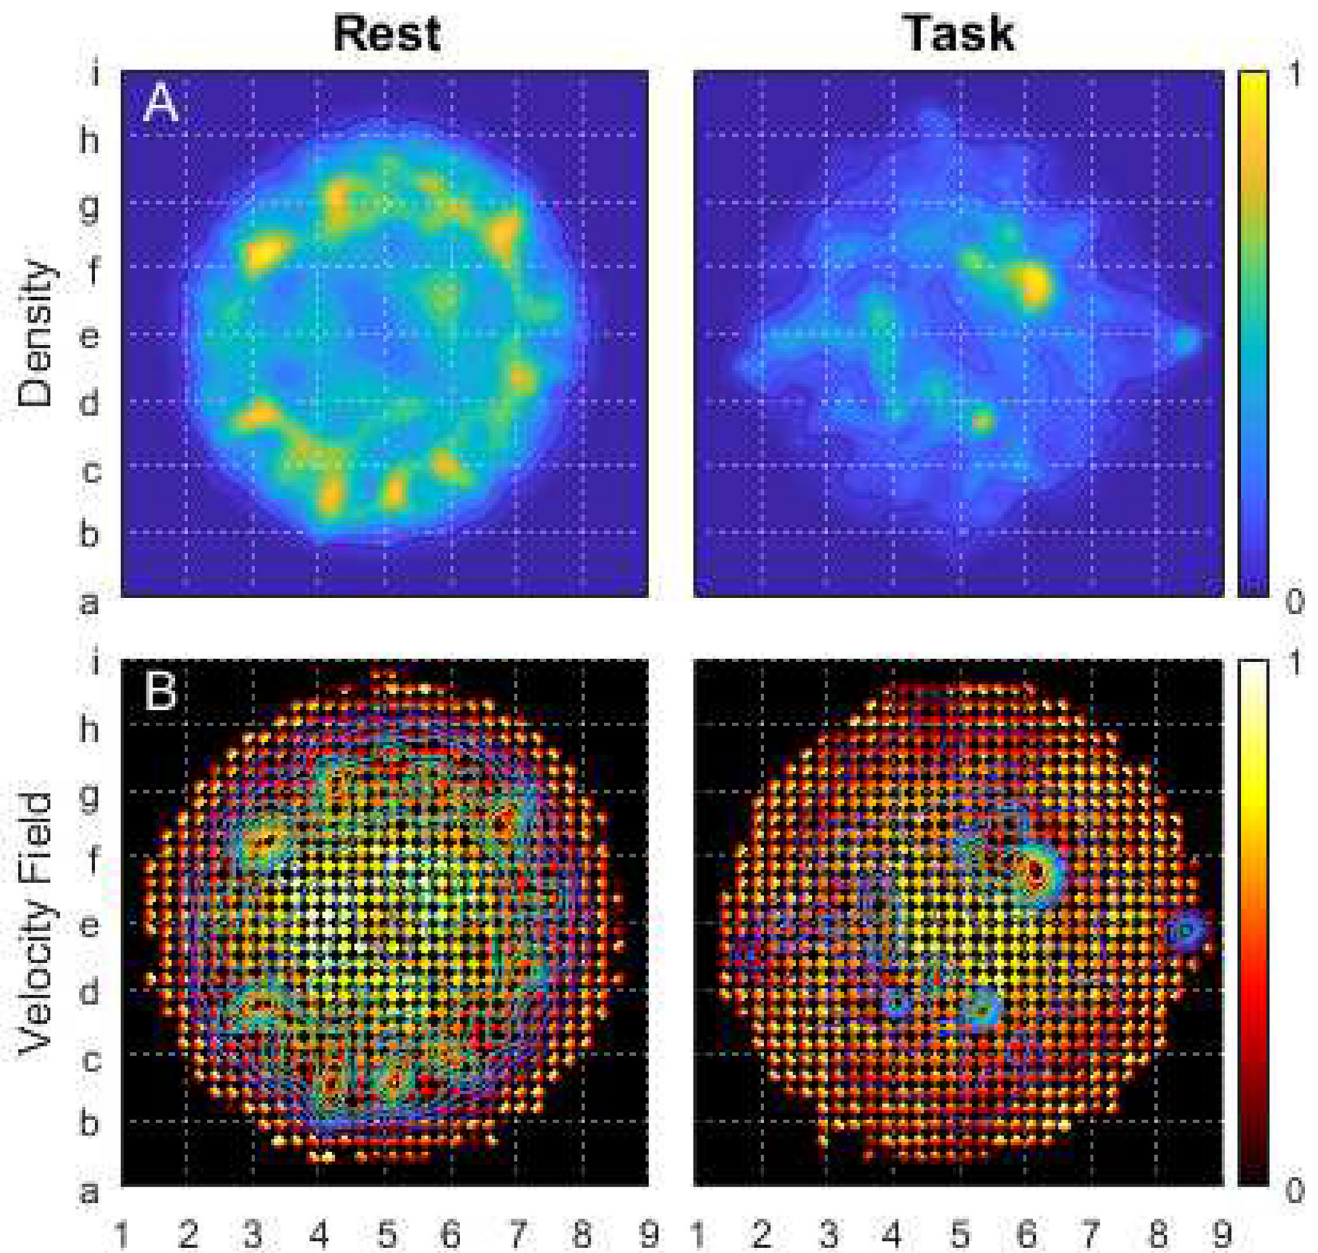


Figure 1.

compares 2D Euclidean embeddings of instantaneous brain states, separated for the resting state and the task-active state. For ease of reference, the embedding space is divided into an 8×8 grid. Alphanumeric labels mark grid vertices. Part A displays the distribution of embedded points as a Gaussian cloud. The Gaussian filter radius equaled 1/32 the maximum point displacement from the embedding's center. Part B displays the velocity field from an aggregation of points within a 32×32 grid in each of the 4 cardinal directions across the map space. All results were normalized to unit magnitude.

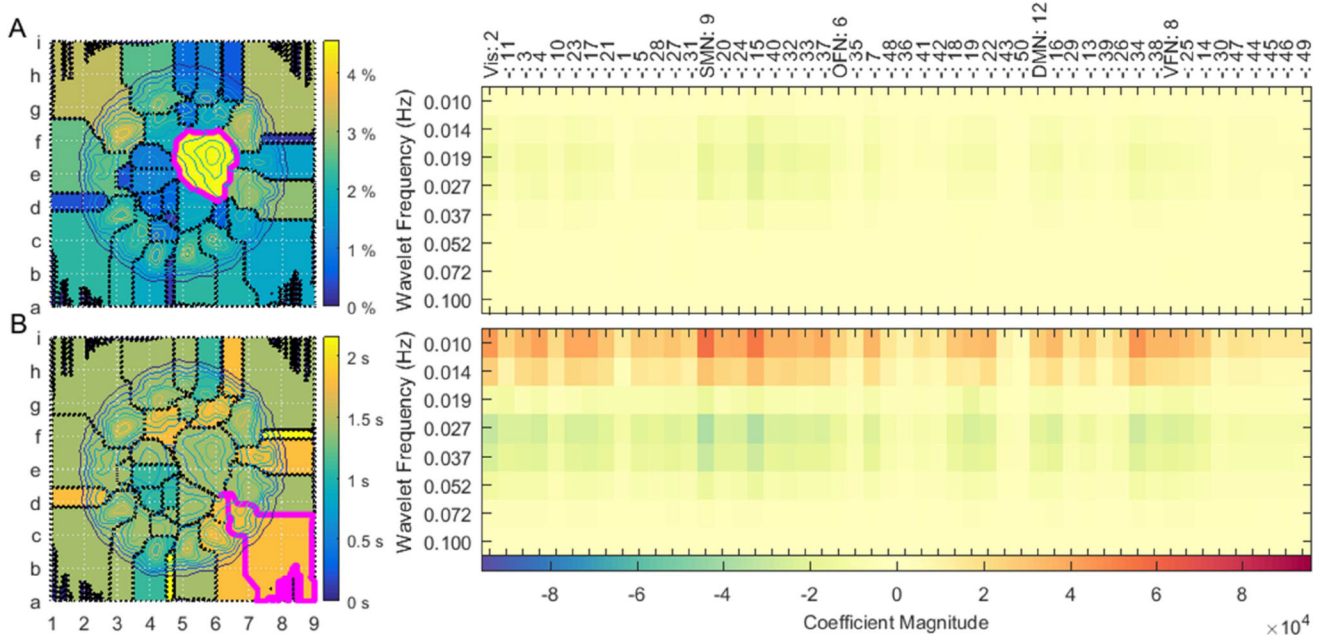


Figure 2. highlights the brain states adopted within watershed regions of the embedding space as participants reside in the resting-state. Part A (left) displays the percentage of points lying within each region. Part B (left) displays the median amount of time participants dwelled in each region. A similar analysis is performed for the task data in supplemental figure S1. The mean spatio-spectral brain state from the regions highlighted in magenta (left) are charted to the right. Each of the 50 ICA resting-state networks are categorized into one of five classes: ‘Vis,’ visual network; ‘SMN,’ somatomotor network; ‘OFN,’ orbito-frontal network; ‘VFN,’ ventral-frontal network.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

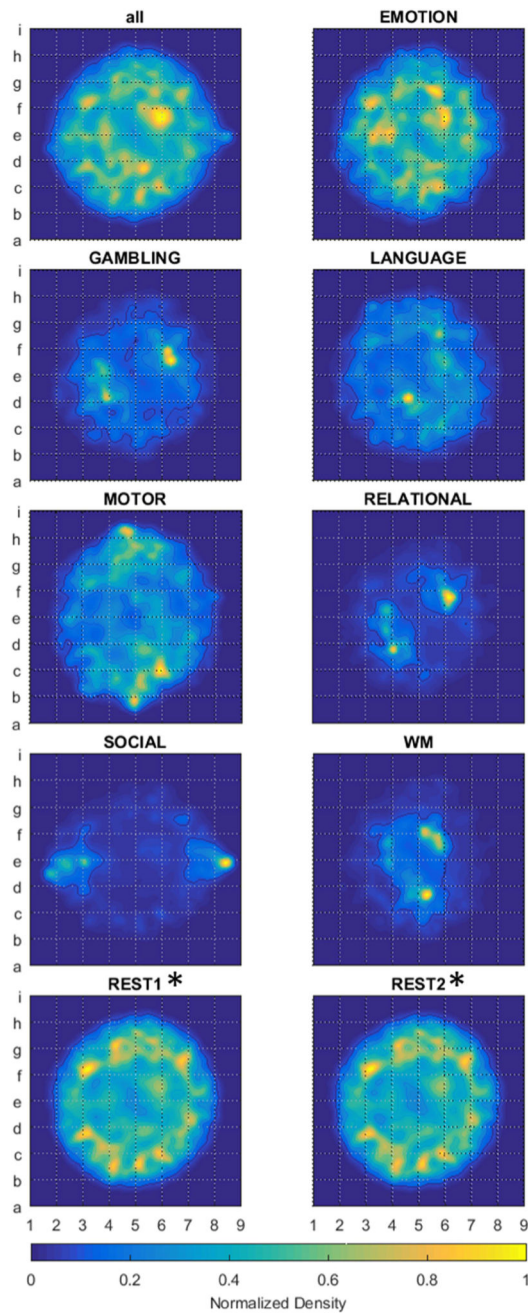


Figure 3. displays t-SNE embeddings of instantaneous brain states, segmented by scan, and represented as the normalized density of each scan's embedded points. Task datasets include BOLD images during all periods of the scan, including any cue events, all contrasting task stimuli, any responses from volunteers, and any fixation blocks. The Gaussian filter radius equaled $1/32$ the maximum point displacement from the embedding's center




Figure 4. plots the mean structural similarity index (SSIM) between normalized density embeddings, segmented across scans. For each scan type, the sample distribution was bootstrapped from 50 realizations of 2500 timepoints, randomly sampled from the group-level data set. The number of time points provides a representative sampling of each scan’s embedded distribution. Asterisks indicate between-scan comparisons whose mean SSIM was not significantly less than either within-scan mean SSIM.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript




Figure 5. analyzes 2D Euclidean embeddings of instantaneous brain states segmented in terms of both block-design task contrasts, and by the acquisition time of temporally aligned task blocks during the SOCIAL scan. The bar chart shows the mean value of the bootstrapped structural similarity index of the within (blue=*mentalizing*, and yellow=*random*) and the between (green) block-design-contrast embeddings. For each task contrast, the sample distribution was bootstrapped from 25 realizations of 250 timepoints, randomly sampled from the group-level data set. Fewer bootstrap time points are used to accommodate the reduced sample size in each segmentation. Daggers above the bar chart point to the aligned block times whose embedding density images are displayed in the inset images. Red boxes define the boundary regions given closer scrutiny in figure 6.

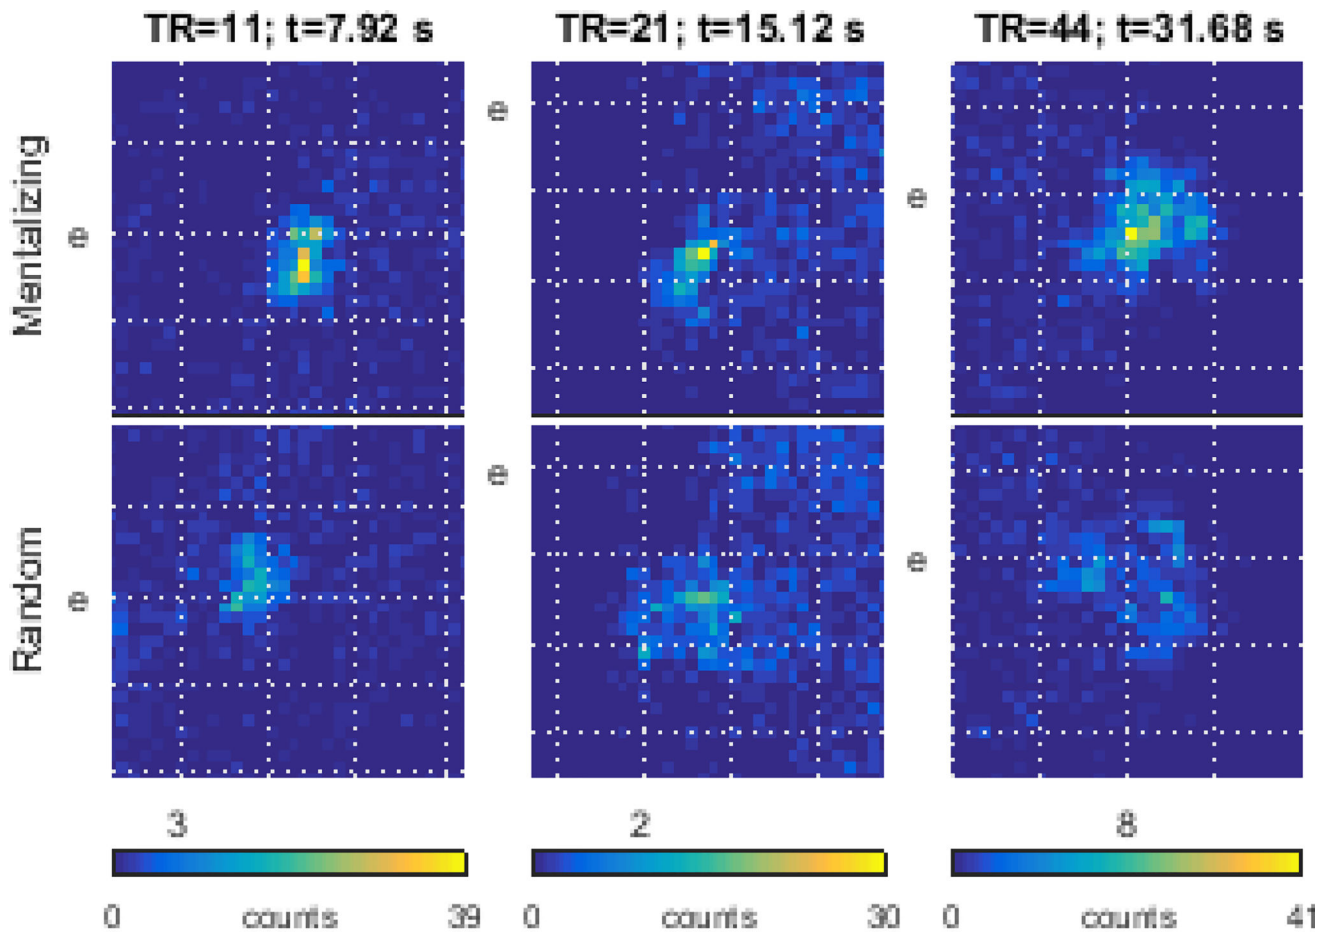


Figure 6. details the brain-state differences between task contrasts by displaying close-up views of the boxed areas from figure 5. Boxed areas are 1/8 the total map space, on a side. Each column is from the same map boundary region. Data are displayed as 2-dimensional histograms from a 32×32 grid inside each box. Column colorbars share the same upper limits.

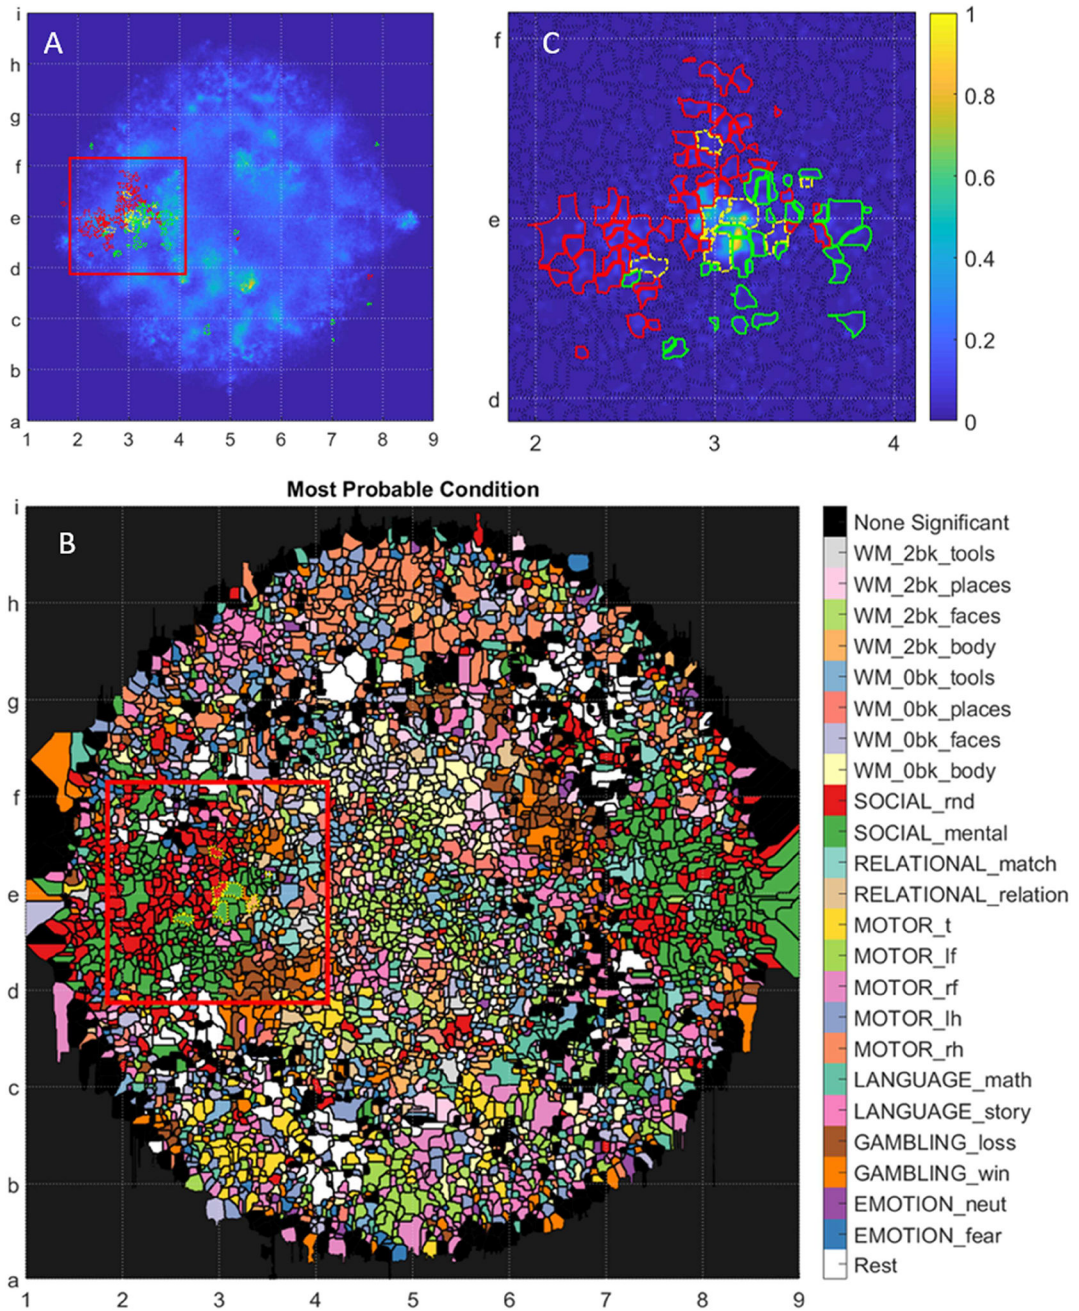


Figure 7. displays the statistical affinity of time-resolved and condition-dependent brain-states for a set of watershed map region. Map regions were segmented using a fine-grained density map generated from all studied time points. This granularity was motivated by the focal organization of map points in figure 6. Here, the filter width was set to $1/256^{\text{th}}$ the distance of the furthest map point from the map's center. Part B displays the most probable state associated with each watershed map region. Regions where no significantly associated state was found were marked in black. Part C pursues the hypothesis that volunteers adopt different brain states around the 11th image of the SOCIAL task when presented with either

the *mentalizing* (green) or the *random* (red) stimulus. Map regions significantly populated in response to either stimuli, at this instant, are outlined in yellow, while regions statistically populated by only one stimulus are outlined in their respective colors. Part C's density map is from only SOCIAL scan data. The red boxes in parts A and B outline the range of part C.

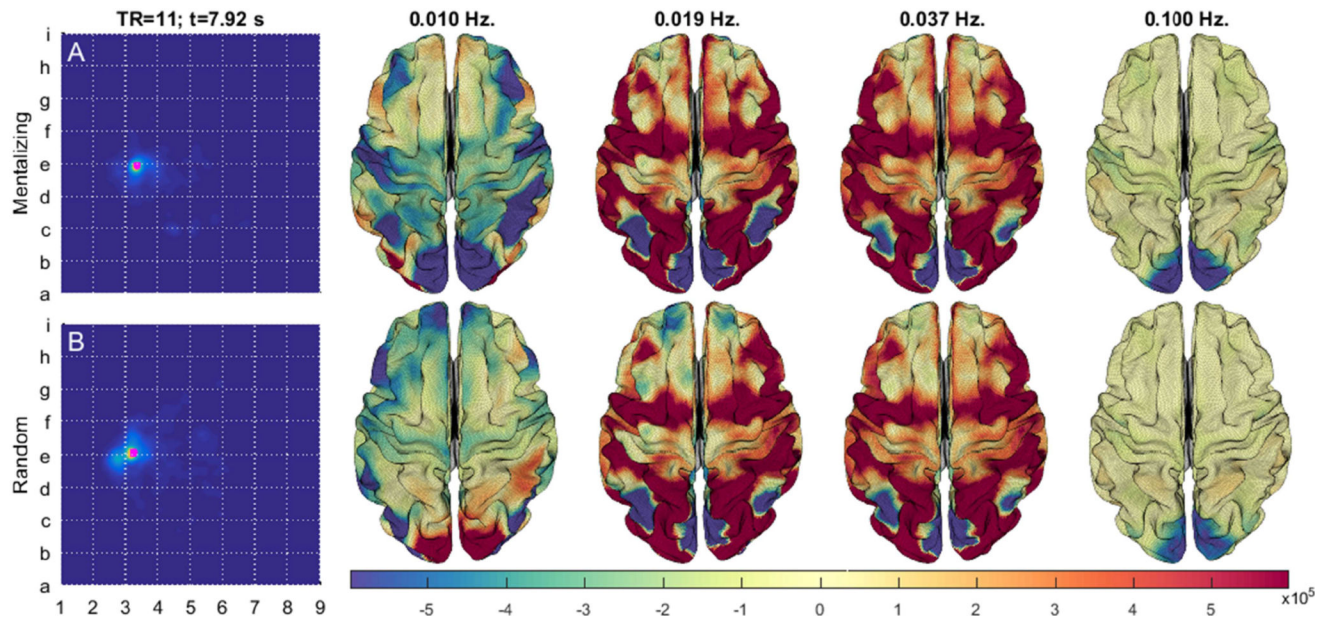


Figure 8. displays the difference between the brain states of participants at ~8 seconds after being given either a *mentalizing* (A) or the *random* (B) visual stimulus during the SOCIAL scan. The brain state is an average from the maximally populated embedding region from within a 32×32 grid (magenta points).