



Latency analysis of resting-state BOLD-fMRI reveals traveling waves in visual cortex linking task-positive and task-negative networks

R. Hindriks^{a,b,c,*}, Mantini R^{f,g}, Gravel N^{h,i}, Deco G^{b,c,d,e}

^a Department of Mathematics, Faculty of Sciences, VU Amsterdam, the Netherlands

^b Center for Brain and Cognition, Computational Neuroscience Group, Department of Information and Communication Technologies, Universitat Pompeu Fabra, Roc Boronat 138, Barcelona, 08018, Spain

^c Institució Catalana de la Recerca i Estudis Avançats (ICREA), Passeig Lluís Companys 23, Barcelona, 08010, Spain

^d Department of Neuropsychology, Max Planck Institute for Human Cognitive and Brain Sciences, 04103, Leipzig, Germany

^e School of Psychological Sciences, Monash University, Melbourne, Clayton VIC, 3800, Australia

^f Research Center for Motor Control and Neuroplasticity, KU Leuven, Tervuursevest 101, 3001, Leuven, Belgium

^g Brain Imaging and Neural Dynamics Research Group, IRCCS San Camillo Hospital, via Alberoni 70, 30126 Venice, Italy

^h Neural Dynamics of Visual Cognition Group, Department of Education and Psychology, Freie Universität Berlin, Berlin, Germany

ⁱ Laboratory of Experimental Ophthalmology, University Medical Center Groningen, University of Groningen, Groningen, The Netherlands

ABSTRACT

Due to the low temporal resolution of BOLD-fMRI, imaging studies on human brain function have almost exclusively focused on instantaneous correlations within the data. Developments in hardware and acquisition protocols, however, are offering data with higher sampling rates that allow investigating the latency structure of BOLD-fMRI data. In this study we describe a method for analyzing the latency structure within BOLD-fMRI data and apply it to resting-state data of 94 participants from the Human Connectome Project. The method shows that task-positive and task-negative networks are integrated through traveling BOLD waves within early visual cortex. The waves are initiated at the periphery of the visual field and propagate towards the fovea. This observation suggests a mechanism for the functional integration of task-positive and task-negative networks, argues for an eccentricity-based view on visual information processing, and contributes to the emerging view that resting-state BOLD-fMRI fluctuations are superpositions of inherently spatiotemporal patterns.

1. Introduction

Over the past two decades, blood-oxygen level-dependent functional magnetic resonance imaging (BOLD-fMRI) has brought about a paradigm shift in our thinking about brain function. BOLD-fMRI studies revealed indeed that in the absence of controlled stimuli or explicit cognitive or perceptual tasks i.e. during the resting-state, brain activity is organized into distributed networks of correlated fluctuations, which correspond to known functional systems (Damoiseaux et al., 2006). Several methods have been proposed to extract functional networks from resting-state fMRI data, among which seed-based functional connectivity (FC) (Biswal et al., 2010) and spatial independent component analysis (spatial ICA) (Beckmann et al., 1457) are the most widely used. A drawback of seed-based FC is that it does not enable disentangling different functional processes that contribute to FC, at least not when these processes overlap in space. The same is true for spatial ICA, which yields the classical resting-state networks (RSN's) under the assumption of no (significant) spatial overlap. As noted in (Smith et al., 2012), the main reason for the widespread use of spatial ICA in analyzing

resting-state fMRI data is not its biological plausibility, but the requirement of spatial ICA algorithms to have many more samples than variables, which is typically true for spatial samples (spatial ICA), but not for temporal samples (temporal ICA). Using a spatial dimension reduction technique, Smith et al. (2012) applied temporal ICA to resting-state fMRI data, which yielded networks that are markedly different from the networks extracted with spatial ICA. In particular, the networks often comprised both activated and deactivated regions and tended to overlap in space. In the current study we essentially extracted the same networks (see Fig. 1). Besides temporal ICA, several other methods have been proposed to extract functional networks (Liu and Duyn, 2013; Yeo et al., 2014; Karahanoglu and Van De Ville, 2015). In these and similar studies, the extracted networks are usually interpreted as fluctuating more or less independently. This is most clearly the case for temporal ICA, which yields networks with temporally independent time-courses (Hyvärinen, 1999).

One of the reasons why functionally distinct networks are often thought of as independent, is because their time-courses often are uncorrelated. In the case of temporal ICA this is by construction, but

* Corresponding author. Department of Mathematics, Faculty of Sciences, VU Amsterdam, the Netherlands.

E-mail address: r.hindriks@vu.nl (R. Hindriks).

<https://doi.org/10.1016/j.neuroimage.2019.06.007>

Received 6 December 2018; Received in revised form 30 May 2019; Accepted 3 June 2019

Available online 13 June 2019

1053-8119/© 2019 Elsevier Inc. All rights reserved.

spatial ICA and other methods typically yield uncorrelated time-courses as well. Within the field of resting-state BOLD-fMRI, the term “correlation” usually refers to correlation at lag zero and it is insufficiently acknowledged that correlations might also exist at non-zero latencies. Within the last decade or so, and using different analysis methods, several studies have reported reproducible latencies between resting-state BOLD-fMRI signals of healthy subjects, (Cecchi et al., 2007; Majeed et al., 2009, 2011; Garg et al., 2011; Mitra et al., 2014; Thompson et al., 2014; Gravel et al., 2017; Griffa et al., 2017; Raatikainen et al., 2017), specific alterations in latency structure under general anesthesia and slow-wave sleep (Wu et al., 2018; Mitra et al., 2015b), and have explored its use as a biomarker for autism spectrum disorders (Mitra et al., 2017) and stroke (Siegel et al., 2016). These and other studies are leading to the view that ongoing BOLD fluctuations are organized as propagating patterns that are characterized by relative latencies and, as such, are intrinsically spatiotemporal. For example (Majeed et al., 2011), analyzed fast (TR = 300 ms) resting-state BOLD-fMRI data from rats and humans and observed quasi-periodically occurring waves of BOLD activity that propagated through different functional systems. Using simultaneous BOLD-fMRI and local field potential recordings in anesthetized rats, Thompson et al. (2014) demonstrated that such quasi-periodic patterns correlate with infra-slow electrical fluctuations. Mitra et al., 2014, 2015a showed that latency patterns can also be extracted from data with longer acquisition times (2–3 s). These studies add to the evidence presented in studies using other measurement modalities, including local field potentials, voltage sensitive dyes, and calcium imaging, that have reported propagating patterns of cortical activity in several species, under a wide range of conditions, and in diverse cortical regions (Golomb and Amitai, 1997; Contreras and Llinas, 2001; Benucci et al., 2007; Rubino et al., 2006; Wu et al., 2008; Lubenov and Siapas, 2009; Mohajerani et al., 2010; Sato et al., 2012; Takahashi et al., 2015; Townsend et al., 2015; Zanos et al., 2015; Zhang and Jacobs, 2015; Vanni et al., 2017; Mitra et al., 2018).

The existence of signal propagation in resting-state BOLD-fMRI data raises the question how propagation patterns relate to known functional networks and zero-lag FC. Mitra et al. (2014) argued that ongoing BOLD fluctuations propagate both within as well as between RSN's, suggesting that RSN's might be embedded in propagation patterns and hence are part of a more comprehensive spatiotemporal organization that integrates the activity within different RSN's. Viewed from a spatiotemporal perspective, zero-lag FC can be interpreted as reflecting the fact that the involved regions are situated at *the same* phase of a propagating pattern, thus leading to simultaneous activation. The absence of zero-lag FC, on the other hand, might reflect true independence or activation at *different* phases of a propagating pattern, thus leading to sequential activation. A direct demonstration of this view has recently been reported (Matsui et al., 2016). Matsui et al. measured hemodynamic signals in anesthetized mice and observed traveling waves of BOLD signal across the neocortex and found a negative correlation between lag and zero-lag FC. They furthermore observed that propagating waves sequentially activated homologue regions of different RSN's, thereby demonstrating that zero-lag FC arises from propagating activity. Such sequential activation of functionally related regions has been observed in several subsequent studies (Ma and Zhang, 2017; Vanni et al., 2017; Mitra et al., 2018). These and other studies argue for a more comprehensive view on resting-state cortical activity, in which BOLD fluctuations within separate functional systems are linked through propagation of BOLD signals.

In this study we propose a method to characterize the latency structure of ongoing BOLD signals about a given functional network, whether defined through spatial (Beckmann et al., 1457) or temporal ICA (Smith et al., 2012) or some other method (Liu and Duyn, 2013; Yeo et al., 2014; Karahanoglu and Van De Ville, 2015). The method provides information on how the BOLD signals behave in the temporal vicinity of the networks'

peak-activation and allows to detect the presence of signal propagation or flow i.e. entanglement between fluctuations in space and time. The method differs from the approaches taken in (Majeed et al., 2009, 2011; Mitra et al., 2014, 2015a; Thompson et al., 2014) in that, instead of first extracting the latency structure (Mitra et al., 2014, 2015a) or quasi-periodic patterns (Majeed et al., 2009, 2011; Thompson et al., 2014) from the data and subsequently relating it to known functional networks, we select a network of interest and directly characterize its flow. It is therefore appropriate in situations in which the dynamics of a given functional network need to be investigated. To test the hypothesis that a given network is embedded in a flow, we describe the construction of appropriate null-data. The null-data has more general applicability and can be used to assess statistical significance of any aspect of propagation, such as the number of latency maps or the occurrence of quasi-periodic patterns (Majeed et al., 2009, 2011; Mitra et al., 2014, 2015a; Thompson et al., 2014). It consists of repeated randomization of the data-set at hand, in such a way as to destroy the time-asymmetry in its cross-covariance structure, while preserving its auto-covariance structure. From the obtained collection of surrogate data-sets, the null-distribution of any desired lag-based index can be approximated and *p*-values can thus be calculated. We emphasize the importance of appropriate null-data and note that earlier studies on propagation phenomena in BOLD-fMRI data have either not directly tested for the presence of propagation or have used inappropriate null-data. Its importance is illustrated by the field of dynamic FC, in which the majority of studies either have not used any statistical testing, or have use inappropriate null-data, with the consequence that the majority of dynamic FC studies contain claims that have not been properly tested for (Hindriks et al., 2016).

We apply the method to functional networks extracted via temporal ICA i.e. temporal functional modes (TFM's) (Smith et al., 2012). To render temporal ICA applicable, we reduce the number of spatial samples by representing the data in the basis of spherical harmonics, defined on the cortical surface. This reduces the number of spatial samples by two orders of magnitude, while retaining the maximal amount of information. We use resting-state BOLD-fMRI data from 94 subjects of the Human Connectome Project (HCP) (Glasser et al., 2013) and use multiple scanning sessions per subject to extract reliable modes and to establish reproducibility of the results. We find nine reliable TFM's, most of which correspond to those reported in (Smith et al., 2012). Most TFM's are bi-modal, in contrast to the classical RSN's extracted via spatial ICA (Smith et al., 2012). The usual interpretation of the amplitudes of a given mode or network is that they reflect the strength with which the mode is expressed at the respective voxels. We illustrate the usefulness of the proposed method by showing that this interpretation is only valid in the absence of flow: In the presence of flow, the amplitudes partly reflect the timing of ongoing BOLD signals from the respective voxels, relative to the mode's peak-activation. Statistical testing demonstrated that several TFM's are embedded in BOLD signal flow. Analysis of one of these TFM's revealed that fluctuations in the task-positive network (TPN) and the task-negative network (TNN) are linked through traveling BOLD waves in early visual cortex (visual areas V1-V4). This observation provides further evidence for the existence of spatiotemporal building blocks of human brain function in which ongoing cortical fluctuations are entangled in space and time (Majeed et al., 2009, 2011; Mitra et al., 2014, 2015a; Thompson et al., 2014). Interestingly, the traveling waves are initiated in the periphery of the visual field and propagate towards the fovea and do this synchronously across hemispheres and visual areas (V1-V4). The observed waves explain why zero-lag FC in BOLD-fMRI within early visual cortex is eccentricity-based (Arcaro et al., 2015) as iso-eccentric locations are activated at the *same* phase of the wave, and why no zero-lag FC is observed between locations with different eccentricities (Arcaro et al., 2015) as such regions are activated at *different* phases of the wave. It thus appears, at least in early visual cortex, that

zero-lag functional connectivity arises from propagating activity, providing more evidence for the embedding hypothesis of cortical activity (Matsui et al., 2016).

2. Materials and methods

2.1. Resting-state BOLD-fMRI data

We provide a short description of the acquisition and pre-processing of the data used in this study. Details can be found in (Glasser et al., 2013; Salimi-Khorshidi et al., 2014). We used the resting-state BOLD-fMRI data from the 94 subjects of the Human Connectome Project that also underwent resting-state magnetoencephalographic (MEG) recordings (Larson-Prior et al., 2013). These subjects were chosen to enable cross-modal comparisons in a subsequent study. Each subject was scanned four times using gradient-echo echo-planar imaging with a 3T Siemens “Connectome Skyra” scanner for 15 min. The subjects were asked to lie still and fixate at a white cross-hair on a dark background, think of nothing in particular, and not to fall asleep. The four scans correspond to two different phase-encoding directions (left-right and right-left), each used two times and were used to verify the reproducibility of the results. This yielded data with a cubic resolution of 2 mm and an acquisition time of 0.72 s. The data were subsequently registered to standard HCP cortical meshes. Nuisance signals such as motion-derived artefacts and physiological noise were cleaned using ICA-FIX, a classifier approach that removes “bad” components from the data. Detailed accounts of ICA-FIX performance can be found in (Salimi-Khorshidi et al., 2014; Griffanti et al., 2014). These procedures yield the FIX-denoised-compact HCP data (Smith et al., 2013) which was used as the starting point for our analysis. Prior to analysis, the BOLD time-series were bandpass filtered between 0.01 and 0.1 Hz using a zero-phase fourth-order Butterworth filter.

2.2. Dimensionality reduction

To reduce the spatial dimension of the data, we made use of the spherical topology of the cortical meshes, by representing the data in the spherical Fourier domain i.e. as linear combinations of (real-valued) spherical harmonics. Specifically, let X_i be the $s \times t$ data matrix of HCP subject i for $i = 1, \dots, 94$, where $s = 64984$ denotes the number of cortical vertices and $t = 1200$ denotes number of temporal samples. We express X in the basis of spherical harmonics up to a given order m_{\max} . The harmonics were sampled on the HCP standard spherical meshes (separately for the left and right hemispheres) and mapped to the convoluted cortex using the homeomorphic spherical registration mapping (Glasser et al., 2013). The homogeneous sampling of the HCP standard cortical sphere ensured that the discretized harmonics are orthogonal. Small distortions in their norms were corrected by dividing every sampled harmonic by its Euclidean norm. The discretized spherical harmonics thus form an orthonormal basis on the left- and right-handed HCP standard cortical meshes.

The use of a maximum order m_{\max} reduces the spatial dimension of the data and also serves to increase its signal-to-noise ratio through its smoothing effect. The number of basis vectors up to and including m_{\max} equals $(2m_{\max} + 1)^2$. The maximum order was set to $m_{\max} = 20$, leading to a $2 \times 441 = 882$ -dimensional basis (441 basis vectors for each hemisphere). This choice still allows spatial detail to be discerned within the regions-of-interest of the cortical parcellation (Glasser et al., 2016b), while significantly increasing the signal-to-noise ratio. The corresponding basis transformation matrices H_L and H_R of the left and right cortical meshes contain these basis vectors as columns and are thus $s \times (2m_{\max} + 1)^2$ -dimensional. The columns of H_L and H_R are sorted in increasing order i.e. decreasing wavelength. The transformation matrix H of the entire cortex (i.e. left and right) is the $s \times 2(2m_{\max} + 1)^2$ which has a block-diagonal structure with H_L and H_R as blocks. All analyses were

performed on the projected matrices $X_i^H = H^T X_i$ and the results were transformed back to the vertex domain via left-multiplication by H . Note that this approach is restricted to analyses that can be implemented by a linear transformation of the data matrix, such as the calculation of covariance matrices or independent component analysis. For group-level analysis, the matrices X_i^H were normalized through entry-wise division by their Frobenius norms.

2.3. Extraction of cortical states

We use the general term *cortical state* to refer to any scalar-valued function defined on the cortex. It can be a (innovative) co-activation pattern (CAP or iCAP) (Liu and Duyn, 2013; Karahanoglu and Van De Ville, 2015), a spatial or temporal independent component (Beckmann et al., 1457; Smith et al., 2012) or any other weight vector obtained from a decomposition or clustering of the data (Yeo et al., 2014). We focus on cortical states that are extracted using temporal ICA i.e. on *temporal functional modes* (TFM's) (Smith et al., 2012). The methods described in Sections 2.4–2.6 are applicable to general cortical states, however. The extraction of TFM's from resting-state BOLD-fMRI data is challenged by the large number of spatial samples as compared to the number of temporal samples. In (Smith et al., 2012), this issue is dealt with by applying temporal ICA to the time-courses of a prior high-dimensional spatial ICA analysis. We deal with it by projecting the data onto the linear space spanned by the cortical spherical harmonics up to a given order (see Section 2.2). This projection reduces the spatial dimension of the data by two orders of magnitude and essentially provides spatial smoothing without serious loss of information.

Temporal ICA was applied to the subject-concatenated and projected data matrix \tilde{X} , which has dimensions $k \times nt$, where $k = 882$, $n = 94$ is the number of subjects, and $t = 1200$ is the number of samples per scan. The matrix \tilde{X} is constructed by concatenating the normalized data matrices \tilde{X}_i of the individual subjects along the temporal dimension. Based on the eigenspectrum of \tilde{X} , the number of independent components was set to 20 and principal component analysis (PCA) was used to project the data matrix onto the 20-dimensional PCA subspace prior to applying temporal ICA. Temporal ICA was applied 20 times and the most consistent decomposition was subsequently selected, where consistency was measured by calculating the maximal Pearson correlation coefficients between the components' time-courses and those of all components of all 19 other decompositions. The correlation coefficients were subsequently averaged over all components. This yielded a decomposition into 20 components for each of the four scanning sessions from which we subsequently selected the most reliable ones. A component was considered reliable if it was present in all four scanning sessions, in the sense of having a (spatial) Pearson correlation coefficient > 0.7 with a (unique) component in each of the three other sessions. This resulted in nine reliable TFM's which are shown in Fig. 1 and will be discussed in Section 3.1.

2.4. Characterization of flow about a cortical state

With every cortical state, one or more time-courses can be associated. For example, if the state is obtained through a spatiotemporal decomposition of the data, its time-course can be taken to be the corresponding temporal component. A time-course of a single data frame, (innovative) co-activation pattern, or cortical state obtained through other methods, can be obtained by calculating a similarity measure between the state and every data frame, the Pearson spatial correlation or covariance being a natural choice. By itself, the time-course associated with a cortical state does not provide information on how the ongoing BOLD fluctuations behave in the temporal vicinity of the state's peak-activations. For instance, the BOLD signals from some cortical region might be activated

simultaneously, relative to the state's peak activation, or their relative timings might be distributed about the state's peak-activation, some peaking earlier, while others peaking later. In the first case, the BOLD fluctuations are not entangled in space and time and their spatiotemporal dynamics can be described by combining a single spatial and temporal function. In the second case, the BOLD fluctuations cannot be separated into a single spatial and temporal function. Rather, the fluctuations are entangled in space and time and *propagate* or *flow*. Below, we describe a method to characterize the spatiotemporal dynamics of ongoing BOLD signals about a given state and to test for the presence of flow.

Let $X : \mathbb{R} \rightarrow \mathbb{R}^n$ be the map that associates with time t the ongoing BOLD activity vector $X(t)$ at time t , where n denotes the number of voxels. Thus, the i -th coordinate X_i corresponds to the BOLD signal at voxel i . Furthermore, let $a : \mathbb{R} \rightarrow \mathbb{R}$ be the time-course of a given brain state $f \in \mathbb{R}^n$. Thus, the cortical state at time t is $a(t)f$. We define the *flow map* $\Lambda : \mathbb{R} \rightarrow \mathbb{R}^n$ of f as the map that associates with lag τ the covariance between X and a at lag τ :

$$\Lambda(\tau) = \int_{-\infty}^{\infty} X(s+\tau)a(s)ds, \quad (1)$$

where the integral is applied coordinate-wise i.e. the i -th coordinate $\Lambda_i(\tau)$ equals the lagged covariance between X_i and a at lag τ . In Eq. (1) we have assumed that X and a are zero-mean stationary signals. The flow map captures the stereotypical BOLD signal as a function of the lag τ , relative to the peak-activations of f , where negative and positive lags correspond to *before* and *after*, respectively. Lagged covariance has been used before in the context of latency analysis of ongoing brain activity under different names such as jitter correlations (Vanni et al., 2017) and temporal-shift map (Wu et al., 2018) and forms the basis of the lag analysis approach described in (Mitra et al., 2014, 2015a). Instead of using it to estimate time-delays as done, for example, in (Mitra et al., 2014, 2015a; Wu et al., 2018), we analyze the map's full algebraic structure.

Specifically, we characterize the latency structure of the flow map by its *flow kernel* $K : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$, which we define as

$$K(\tau_1, \tau_2) = \langle \Lambda(\tau_1), \Lambda(\tau_2) \rangle, \quad (2)$$

where τ_1 and τ_2 are arbitrary lags and the brackets denote taking the inner product. Note that K is symmetric i.e. $K(\tau_2, \tau_1) = K(\tau_1, \tau_2)$ for all $\tau_1, \tau_2 \in \mathbb{R}$. The flow kernel measures the resemblance between the spatial vectors $\Lambda(\tau_1)$ and $\Lambda(\tau_2)$ and can also be viewed as a linear operator that maps a signal x onto another latency signal Kx by the integral transformation

$$(Kx)(\tau) = \int_{-\infty}^{\infty} K(s, \tau)x(s)ds. \quad (3)$$

A signal $x : \mathbb{R} \rightarrow \mathbb{R}$ that associates a value $x(\tau)$ to every lag τ will be referred to as a *latency signal*. Examples of latency signals are the coordinate functions Λ_i of the flow map Λ . Note that the collection of all (square integrable) latency signals forms a vector-space over \mathbb{R} . Thus, the latency structure within the data, relative to the brain state f , is characterized by an integral self-operator on latency space with kernel K . The ongoing BOLD signals about f are considered to *flow* or *propagate* if the columns of the flow kernel K are lagged with respect to each other. Since the columns of K typically peak about zero (see Section 3.2), we can measure the flow strength by quantifying the dispersion of peaks, for example by calculating their standard-deviation. This leads to a *flow index* $\rho \geq 0$ with the dimension of time.

Depending on the research question, the global signal might be removed from the data and it might therefore be useful to be able to relate the flow map and its associated kernel of the transformed data to those of the untransformed data. These relationships can be derived for the most common ways of dealing with the global signal, which are regression, subtraction, and normalization (Liu et al., 2017). In the case

of regression, the transformed data are of the form $X_{\text{GSR}}(t) = X(t) - g(t)\alpha$, where $g(t)$ denotes the global signal and α denotes the spatial vector of regression coefficients. From the definition of the flow map, it directly follows that

$$\Lambda_{\text{GSR}} = \Lambda - \Lambda_g, \quad (4)$$

where

$$\Lambda_g(\tau) = \alpha \int_{-\infty}^{\infty} g(s)a(s+\tau)ds, \quad (5)$$

is the flow map of the global signal. Note that if the global signal is uncorrelated with the state's time-course a for every lag τ , then $\Lambda_g = 0$ and hence $\Lambda_{\text{GSR}} = \Lambda$. Below, we write $\Lambda_g(\tau) = \alpha\gamma(\tau)$. It is straightforward to show that

$$K_{\text{GSR}} = K + K_g + K_\gamma, \quad (6)$$

where $K_g(\tau_1, \tau_2) = \Lambda_g(\tau_1)^T \Lambda_g(\tau_2)$ is the flow kernel of the global signal and

$$K_\gamma(\tau_1, \tau_2) = -[\gamma(\tau_1)\Lambda^T(\tau_1) + \gamma(\tau_2)\Lambda^T(\tau_2)]\alpha \quad (7)$$

is a kernel related to the interaction between the global signal and the state's signal and vanishes in the absence of correlations at all lags. Similar relationships for global signal subtraction and normalization can be derived.

2.5. Detection of signal flow in experimental data

Since experimental data comprise a finite number, say k , samples, the BOLD activity $X \in \mathbb{R}^{n \times k}$ becomes an $n \times k$ (space \times time) matrix, the state's time-course $a \in \mathbb{R}^k$ becomes a (temporal) vector, and the flow map Λ becomes a $n \times (2\tau_{\text{max}} + 1)$ matrix, where τ_{max} denotes the maximal latency and is a free parameter that can be chosen on the basis of the auto-covariance function of a . Furthermore, the kernel K becomes a $(2\tau_{\text{max}} + 1) \times (2\tau_{\text{max}} + 1)$ matrix. A natural estimate of Λ is

$$\hat{\Lambda} = XA, \quad (8)$$

where A denotes the $k \times (2\tau_{\text{max}} + 1)$ matrix given by

$$A = \begin{bmatrix} a_{-\tau_{\text{max}}} & a_{-\tau_{\text{max}}+1} & \cdots & a_{\tau_{\text{max}}-1} & a_{\tau_{\text{max}}} \end{bmatrix}, \quad (9)$$

where a_j denotes a translated in time over j samples. Note that the i -row of $\hat{\Lambda}$ is the sample cross-covariance function between the BOLD signal at the i -th voxel and a . The flow kernel is estimated by

$$\hat{K} = \hat{\Lambda}^T \hat{\Lambda}, \quad (10)$$

which is a $(2\tau_{\text{max}} + 1) \times (2\tau_{\text{max}} + 1)$ matrix. A natural estimate of the flow index, as defined in Section 2.4, is given by the sample standard-deviation of the peak-locations of the columns of \hat{K} . This yields an estimated value $\hat{\rho}$ of ρ which is a natural estimate of the strength of flow about the brain state f . Detection of flow in a given data-set now amounts to testing the null-hypothesis of no flow using the test-statistic $\hat{\rho}$, for which the null-distribution, that is, the sampling distribution of $\hat{\rho}$ under the null-hypothesis, is required. Since this distribution is unknown, it needs to be approximated and we will do this by appropriately constructed surrogate-data. (Schreiber and Schmitz, 2000). Thus, we generate a large number of copies X_{surr} of the data matrix X that have the same statistical properties of X , but lack the property of interest (i.e. flow), and subsequently compute $\hat{\rho}$ from each of them to construct an approximation to the null-distribution. We generate X_{surr} by coherently phase-randomizing the symmetrized cross-spectral matrix of X . This

leaves the auto-covariance structure of X intact but symmetrizes its cross-covariance structure. To generate a surrogate copy of a $t \times n$ data-matrix X , where n denotes location and t time, we first compute its discrete Fourier transform $\tilde{X} = FX$, where F is the $t \times t$ discrete Fourier transformation matrix, and subsequently symmetrize it by taking its real part:

$$\tilde{X}_{\text{symm}} = \text{Re}[\tilde{X}]. \quad (11)$$

We then generate a complex-valued t -dimensional phase-vector $\exp(i\phi)$, where ϕ is a t -dimensional conjugate-symmetric stochastic vector, whose entries are uniformly distributed in the interval $[0, 2\pi]$ and subsequently form

$$\tilde{X}_{\text{symm}} = D_\phi \text{Re}[\tilde{X}], \quad (12)$$

where D_ϕ denotes the $t \times t$ diagonal matrix with $\exp(i\phi)$ on the diagonal. The randomization is coherent in the sense that each row of $\text{Re}[\tilde{X}]$ is multiplied by *the same* complex-valued phase and, as a consequence, the instantaneous covariance matrix of X is preserved (Prichard, 1994). The surrogate copy of X is now obtained by taking the inverse discrete Fourier transform of \tilde{X}_{symm} :

$$X_{\text{surr}} = F^{-1} \tilde{X}_{\text{symm}} = F^{-1} D_\phi \text{Re}[FX]. \quad (13)$$

The surrogate copy X_{surr} has (approximately) the same empirical auto-covariance functions as X , with the difference that they are symmetric, which is due to the symmetrization of the Fourier transform of X .

Since phase-randomization generally does not preserve the amplitude distribution of the data and preserves autocorrelation functions also approximately, we checked the results using a randomization scheme that does preserve these properties. To generate a surrogate copy of the subject-concatenated data matrix X , the data matrix X_i of subject i is reversed in time with probability $p = 0.5$ and this is done independently for each subject. The limitation of this method is that unlike phase-randomization, it can be applied only on the group-level. Furthermore, to obtain an accurate approximation of the null-distribution, a large number subjects is required. The advantage, however, is that no assumptions on the distribution of the data are made.

2.6. Lagged connectivity between functional temporal modes

We consider again the resting-state BOLD-fMRI vector $X(t)$ and suppose that the assumption underlying temporal ICA is satisfied i.e. $X(t)$ is a superposition of spatial components (the TFM's) f_m with temporally independent time-courses $a_m(t)$:

$$X(t) = \sum_m a_m(t) f_m, \quad (14)$$

where, and without loss of generality, we assume the time-courses to have unit variance. In the context of temporal ICA, temporal independence refers to the factorization of the joint probability density of the time-courses at lag zero. It thus implies vanishing correlations between all time-courses at lag zero, but not necessarily at non-zero lags. The flow map $\Lambda_j(\tau)$ of f_j can hence be written as

$$\Lambda_j(\tau) = \gamma_j(\tau) f_j + \sum_{m \neq j} \gamma_{jm}(\tau) f_m, \quad (15)$$

where $\gamma_{jm}(\tau) = \mathbb{E}[a_j(t + \tau) a_m(t)]$ denotes the Pearson correlation between $a_i(t)$ and $a_j(t)$ at lag τ and $\gamma_j = \gamma_{jj}$. Furthermore, the flow kernel, viewed as an operator in lag space, takes the form

$$(Kx)(\tau) = \mu_j \gamma_j(\tau) + \sum_{n \neq j} \mu_n \gamma_{jn}(\tau), \quad (16)$$

where $\mu_j = \mu_{jj}$ and

$$\mu_{jn} = \sum_m F_n^T f_m \left(\int \gamma_{jm}(s) x(s) ds \right). \quad (17)$$

Note that since the time-courses are independent, the flow map at lag zero equals f_j and we can therefore speak of the flow *about* f_j . The flow map can thus be viewed as the sum of two parts. The first part is the contribution of f_j and equals the autocorrelation function $\gamma_j(\tau)$. The second part contains the contributions from all other TFM's and equal the cross-correlation functions $\gamma_{jm}(\tau)$. Note that in the absence of flow Λ_j reduces to $\Lambda_j(\tau) = \gamma_j(\tau) f_j$ so that the flow kernel K_j factorizes as a product of covariances:

$$K_j(\tau_1, \tau_2) = \|f_j\|^2 \gamma_j(\tau_1) \gamma_j(\tau_2). \quad (18)$$

In the presence of flow, $\Lambda_j(\tau)$ has rank > 1 , which implies that $\gamma_{jm}(\tau) / = 0$ for at least one mode f_m with $m \neq k$ and at least one non-zero value of τ . The presence of flow thus implies the existence of correlations between the time-courses of (some of the) TFM's at (some) non-zero lags. Moreover, the flow map can be decomposed into the contributions from other TFM's and depends on the lagged correlation structure between the TFM's. The lagged correlation structure therefore provides an alternative way of analyzing the flow about TFM's. It will be convenient to quantify the contribution of one TFM to the flow about another TFM in such a way that it is independent of lag. We thus define the *lag-averaged correlation* $\bar{\gamma}_{jm}$ as

$$\bar{\gamma}_{jm} = \frac{1}{2\tau_{\text{max}}} \int_{-\tau_{\text{max}}}^{+\tau_{\text{max}}} |\gamma_{jm}(\tau)| d\tau, \quad (19)$$

where the vertical bars denote taking absolute value and where τ_{max} denotes the maximal lag of the flow map (see Section 2.5). The lag-averaged correlation can be estimated from experimental data by replacing the covariance functions γ_{km} by their sample estimates and the integral by a sum.

3. Results

3.1. Temporal functional modes of cortical resting-state BOLD activity

Fig. 1 shows the nine reliable temporal functional modes (TFM's) i.e. the modes that were present in each of the four scanning sessions (see Section 2.3 for details on their extraction). The coefficient of variation of their average amplitudes is 0.39 which means that the modes are comparable in strength. All 20 TFM's extracted from each of the four scanning sessions are displayed in Supplementary Figs. 1–4. Two features are of interest. First, and in contrast to the classical resting-state networks (RSN's), most modes have both positive and negative amplitudes, in line with earlier studies on TFM's (Smith et al., 2012) and co-activation patterns of resting-state BOLD fluctuations (Liu and Duyn, 2013). Following (Smith et al., 2012), we will refer to these as *activations* and *deactivations*, respectively, and to such modes as *bi-modal*. Note, however, that the sign convention used here is arbitrary and that the important feature of the modes is their bi-modality (Smith et al., 2012; Liu and Duyn, 2013). We note that uni-modality of the classical RSN's is due to the spatial orthogonality constraint of the independent components (when using spatial ICA) and the aggregation of distinct functional processes (when using seed-based FC), which complicates the functional interpretation of these networks (Smith et al., 2012). Second, a single TFM receives contributions from functional nodes of different RSN's and functional nodes of a single RSN can be observed in several TFM's. These features reproduce earlier reports on TFM's and other functional modes (Smith et al., 2012; Liu and Duyn, 2013) and suggest a functional organization of the resting human cortex that is fundamentally different from that suggested by the classical notion of RSN's.

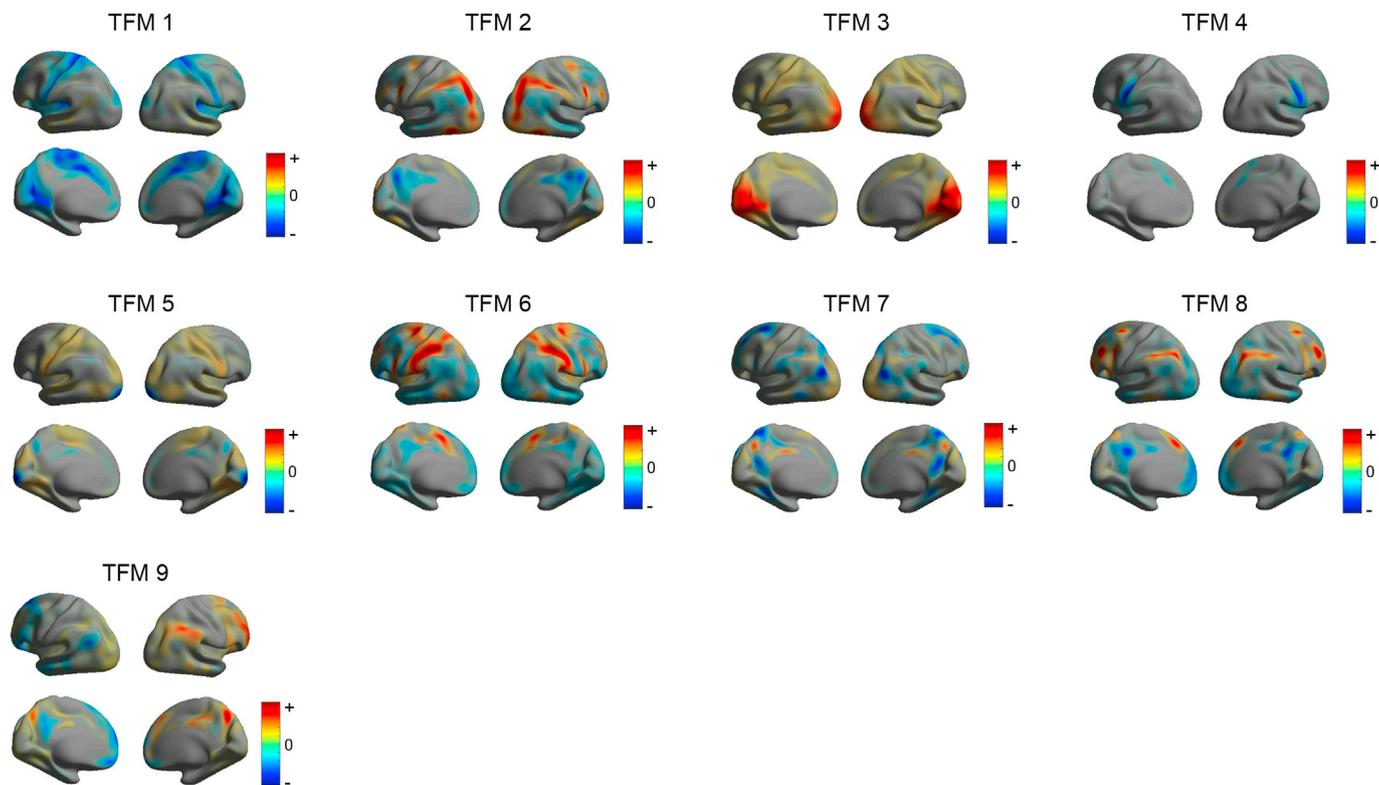


Fig. 1. Temporal functional modes of cortical resting-state BOLD-fMRI activity. A. Temporal functional modes (TFM's) extracted on the group-level ($n = 94$) using temporal independent component analysis (temporal ICA). All modes were observed in each of the four scanning sessions. Shown are the session-averaged modes. Each mode has been scaled to fill the full colorbar. B. Amplitudes of the modes displayed in A. as measured by the standard deviation of their spatial profiles (their temporal profiles are standardized to standard deviation one).

3.2. BOLD signal flow through temporal functional modes

To detect the presence of flow about the extracted TFM's, we computed their flow indices on the group-level from the $882 \times nt$ -dimensional subject-concatenated and projected data matrix \tilde{X} . The lags ranged between $\tau = \pm 19$ samples, which corresponds to about ± 13.7 seconds ($TR = 0.72$ s). The flow maps were computed for every subject separately and subsequently averaged. Fig. 2A shows the obtained index values for each of the modes and for all four scanning sessions. The figure shows that the flow indices are reasonably reliable across sessions. This is especially true for Mode 5 (sensorimotor-DMN), the flow about which is by far the strongest and the differences in flow strength across sessions are about 10% of its average value. Statistical significance was assessed for each session separately, by comparing each of the nine observed index values to the 95%-percentile of their distribution under the null-hypothesis of no flow. The null-distributions were approximated by using a large number of phase-randomized surrogate copies of the subject-concatenated data matrix (see Section 2.5), extracting 20 TFM's from each copy, and subsequently calculating the flow indices for each mode. Thus, all nine observed index values were compared against the same null-distribution. In constructing the distribution, only the maximal index value of each surrogate copy was kept. This circumvented the complication of having to pair the modes extracted from the observed and surrogate data. A consequence of using the maximal values is that the resulting statistical test tends to be conservative, because the observed index value of each of the modes is compared against the maximal index value over all 20 modes extracted from the surrogate data. A mode was considered to flow if the mode's index exceeded the 95% percentile of the approximated null-distribution, after correcting for multiple (9) comparisons using Bonferroni. With 179 surrogate copies, this implies significant flow with a corrected p -value < 0.05 , if the observed index value

exceeds all surrogate values. Applying the test to all four scanning sessions showed that all non-zero values in Fig. 2A are statistically significant. We repeated the testing procedure using the surrogate data obtained by random time-reversal of the subject's data matrices (see Section 2.5) which showed that only the flow about TFM 5 is significant (in all four scanning sessions). This discrepancy is to be expected as the surrogate data obtained by random time-reversal corresponds to a slightly more general null-hypothesis and hence is harder to reject.

Fig. 2B shows the flow kernels of all nine TFM's, computed from scanning session RL₁. Its value at a given pair of lags (τ_i, τ_j) corresponds to the spatial covariance between the i -th and the j -th frames (i.e. columns) of the mode's flow map. The kernels are symmetric by construction. The figure shows that the columns of the flow kernel have a single peak. The flow index measures the dispersion of the peak-locations, which is particularly large for Mode 5. The dispersion of peak-locations is more apparent in Fig. 2C, which shows the flow kernel columns. The columns are plotted for lags within ± 8 seconds to more clearly show the peak-locations. We see that for Modes 6 and 7, all maxima are located at lag zero and hence their observed index values are zero (see Fig. 2A). We also see that the columns of Mode 5 peak at different latencies, which is reflected in a large observed index value.

3.3. Latency effects on functional mode amplitudes

Most TFM's are bi-modal i.e. they comprise activated as well as deactivated regions (see Section 3.1 and Fig. 1), which is consistent with earlier observations on TFM's and other kinds of functional modes (Smith et al., 2012; Liu and Duyn, 2013). In (Smith et al., 2012; Liu and Duyn, 2013), bi-modality was interpreted as reflecting anti-correlations between BOLD signals from different locations, an interpretation that is consistent with other studies on anti-correlations in resting-state

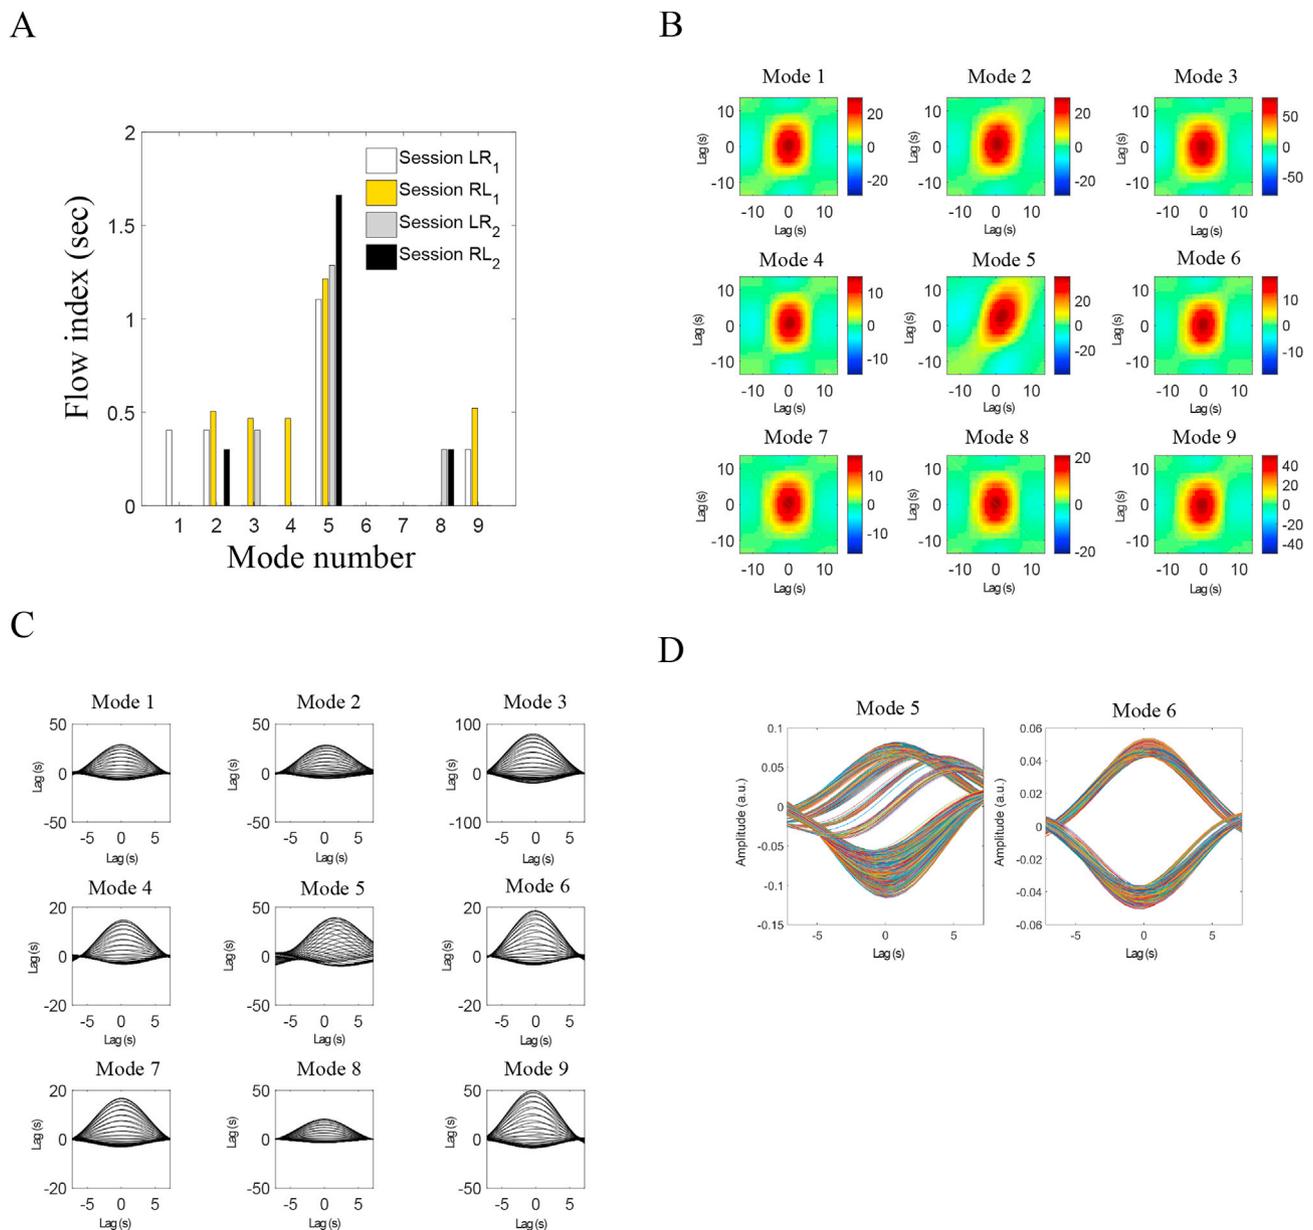


Fig. 2. Flow about temporal functional modes. **A.** Observed values of the flow index for each of the nine temporal functional modes and for each of the four scanning sessions. **B.** Flow kernels of all nine modes. **C.** Column-wise representation on the flow kernels shown in panel B. The range of lags is $\tau = \pm 8$ seconds for better visibility. **D.** Shown are the 500 latency signals (i.e. rows of the flow map) with the highest variance for Mode 5 (left panel) and Mode 6 (right panel). The latency signals were selected separately for cortical locations that contribute negatively/positively to the mode. Thus, in total, 10^3 signals are displayed.

BOLD-fMRI (Fox et al., 2005, 2009; Fransson, 2005). The amplitude of a voxel within a functional mode can then be interpreted as the voxels' contribution to the mode or, equivalently, as the expression-strength of the mode at that voxel. This interpretation, however, only holds when BOLD fluctuations about the mode are indeed anti-correlated. In the presence of flow, however, voxel amplitudes are also influenced by the *timing* of BOLD fluctuations about the mode. Bi-modality of a TFM is generally insufficient to conclude that BOLD fluctuations in activated and deactivated regions are anti-correlated, except in the special case of no flow. To establish anti-correlations about a mode, the spatiotemporal dynamics of BOLD fluctuations in a temporal window about the mode need to be taken into account. This information is contained in the flow map at non-zero lags τ .

Compare, for instance, TFM 6 (DAN), which has flow index zero, and TFM 5 (sensorimotor-DMN), which has a high (and significant) flow index. Fig. 1 clearly shows the bi-modality of TFM 6: regions inside and

outside the DAN are activated and deactivated, respectively. Since a flow index of zero is equivalent to anti-correlations of BOLD signals about the mode, a zero flow index suggests that the ongoing BOLD fluctuations about the mode are anti-correlated. This can in fact already be concluded from its flow kernel shown in Fig. 2B and C. The anti-correlations can be more directly observed in Fig. 2D (right panel) which shows the 10^3 latency signals (i.e. rows of the flow map) with the highest variance, selected separately for vertices with positive and negative amplitudes. Each curve corresponds to a cortical location and describes the stereotypical behavior of the BOLD signal at that vertex about the mode's peak-activation (which is at lag zero). The figure shows that the BOLD fluctuations about TFM 6 indeed are anti-correlated and, as a consequence, the amplitudes of TFM 6 can be interpreted as contributions as done in (Smith et al., 2012; Liu and Duyn, 2013). This interpretation is valid for all TFM's with zero flow index. TFM 5, in contrast, has a high flow index and inspection of its flow kernel shows that BOLD fluctuations about this

mode are not anti-correlated, but instead, peak at distinct lags (see Fig. 2B and C). This can be more directly observed in Fig. 2D (left panel), which shows that ongoing BOLD signals can peak both *before* as well as *after* the mode itself, with latencies up to about 5 s. As a consequence, the amplitudes of TFM 5 cannot be interpreted as expression-strengths because they also depend on the *timing* of the BOLD signals, relative to the mode's peak-activation. For example, the BOLD signals that peak several seconds after the mode itself, contribute negatively to the mode at lag zero.

3.4. Coordination between task-negative and task-positive networks

In the previous section we showed the existence of BOLD signal flow about some of the TFM's and discussed the implications for the interpretation of TFM amplitudes. Although flow kernels provide information about the relative timing of BOLD signals, they do not provide information about the flow's spatiotemporal organization. In this section we consider the organization of BOLD flows in space and time by analyzing the frames (i.e. columns) of the corresponding flow maps. Because the latency signals (i.e. rows of the flow map) of non-flowing modes peak at the same time, these modes behave like standing waves i.e. are not entangled in space and time and therefore have a trivial spatiotemporal structure. We therefore focus on TFM 5, which has the highest flow index of all TFM's. We will analyze the flow map extracted from scanning session RL₁. The flow maps of the other three scanning sessions are similar and are shown in Supplementary Figs. 5–7. Moreover, Supplementary Fig. 8 shows the flow map extracted from the data of a single

subject and using a higher value of the maximal order of the Harmonic basis ($m_{\max} = 40$) which corresponds to less spatial smoothing (see Section 2.2). The figure demonstrates that BOLD flow can be observed in single-subject data and is robust to the precise value of m_{\max} .

Fig. 3A shows the frames of the flow map of TFM 5, represented as a series of cortical flat maps. The flow map is displayed within the range $\tau = \pm 14$ samples and lags are relative to the mode's peak-activation at lag zero. Note that the flow is highly synchronized across hemispheres (average spatial correlation between left- and right-hemispheric frames is 0.94), most likely due to either inter-hemispheric anatomical connectivity or coordinated thalamic drive. To highlight the spatial pattern within individual frames, the colors of each frame were not scaled relative to the other frames so that amplitude differences between frames cannot be assessed. Fig. 3B shows the same frames but now scaled so that amplitude differences can be compared. The drop-off in amplitude for large (absolute) lags reflects decreased lagged correlations between the ongoing BOLD signals and the time-series of the functional state.

To analyze the flow in terms of functional anatomy, consider Fig. 3C and D, which show the propagation map of TFM 5 at lag zero on flat and inflated cortical maps, respectively. The most strongly contributing regions are labeled. Negatively contributing vertices cover the foveal section of early visual cortex, together with task-negative regions, namely (parts of the) posterior and anterior cingulate cortices, temporal lobe, supramarginal and angular gyri, insular cortex, as well as large parts of the frontal lobe. The contributions of regions other than the fovea and the PCC (Area 23d) are relatively weak, but nevertheless observable in all four scanning sessions (see SF 6-8), and hence robust. We identify this set

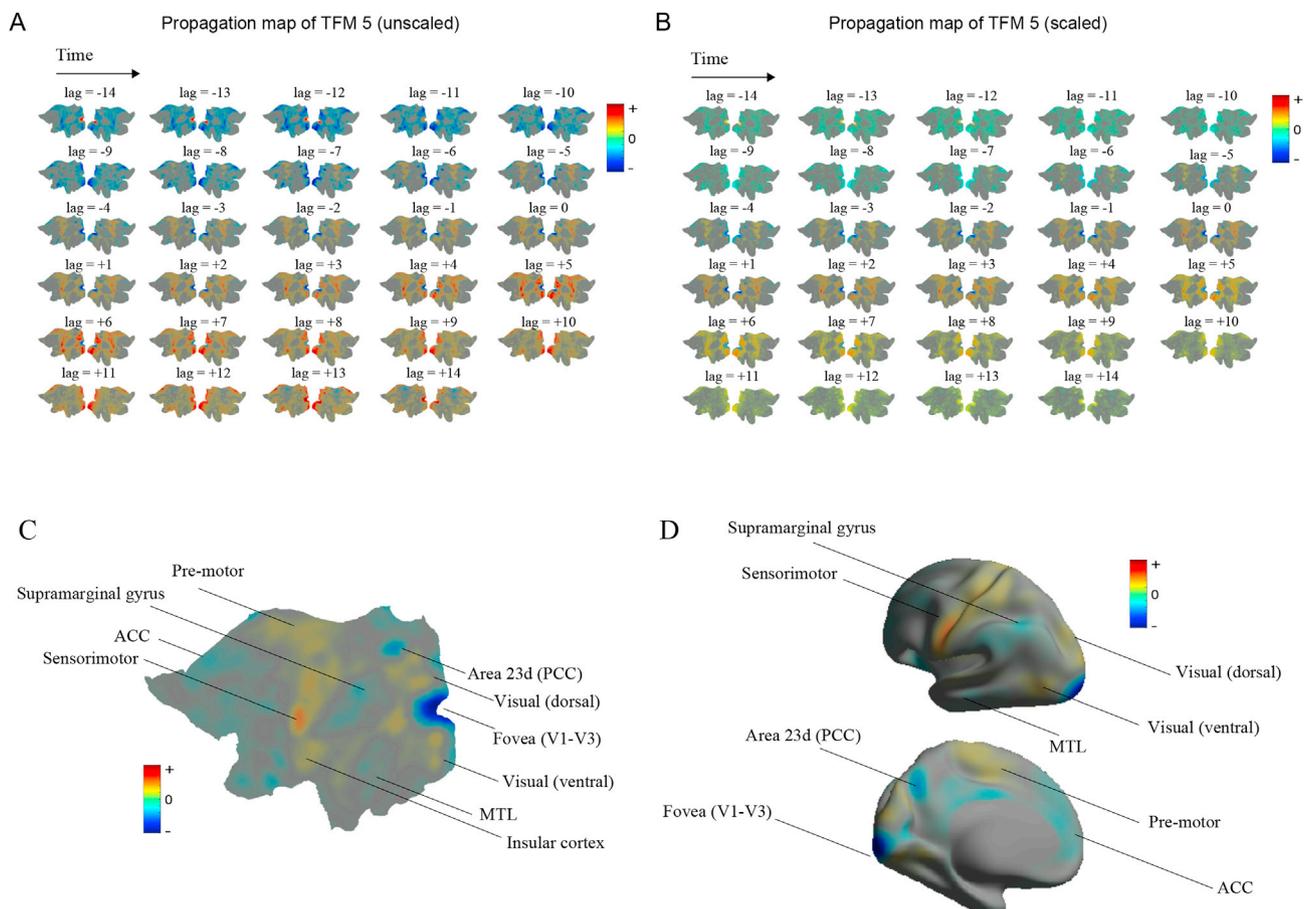


Fig. 3. BOLD signal flow about TFM 5. A. Visualization of the flow map of TFM 5 (sensorimotor-DMN) as a series of frames, displayed on cortical flat maps. Lags range from $\tau = -14$ samples to $\tau = +14$ samples, relative to the mode's peak-activation at $\tau = 0$. Colormaps are unscaled so that amplitudes cannot be compared across frames. Time is in the reading direction. B. Same as in A. but with scaled colorbars so that amplitudes can directly be compared across frames. C. Mode 5 displayed on a cortical flat map (left hemisphere only) in which the most important contributing regions are labeled using (Glasser et al., 2016a). D. Same as in panel C but using an inflated cortical mesh for easier recognition of the functional regions.

of regions with the task-negative network (TNN). The second network essentially corresponds to the complement of the TNN and covers the peripheral parts of early visual cortex, higher-order visual cortex (both dorsal and ventral), primary- and higher-order somatosensory cortex, as well as primary and pre-motor cortex. These regions contribute to TFM 5 in all four scanning sessions (see SF 6–8) and hence they are robust as well. We identify this set of regions with the task-positive network (TPN). The TNN and TPN can also be separated by using the relative timing of the lag signals. In (Vanni et al., 2017), the topography of oscillatory activity was characterized by plotting the phases of the second coefficient of the Fourier transform as a function of location. These coefficients (i.e. the first harmonics) reflect the relative timing of the slowest oscillations within the data. Inspection of the Fourier spectra of the lag signals of TFM 5 showed that they mainly differed in the phases of the second and third harmonics. We thus plotted the difference in phase between the second and third harmonic as a function of location. The resulting topography is shown in Fig. 4A. A comparison with Fig. 3C and D shows that the positive and negative phase-differences correspond with TPN and TNN, respectively.

To see how the dynamics about this mode relates to (de)activation of the TPN and TNN, in Fig. 4B we plotted the average latency signals within TPN and TNN. We see that TPN and TNN, respectively, reach their peak (de)activation about 2 s before and 2 s after the mode's peak-activation. Going back to Fig. 3, we see that BOLD signals flow both *within* TPN and TNN and *between* TPN and TNN. Consider, for example, the TPN at $\tau = -14$ samples, at which the sensorimotor cortex is deactivated over its entire lateral extent and we can observe a localized deactivation within the insular cortex. At about $\tau = -7$, the sensorimotor cortex gets activated, while the insular region is still deactivated. This implies the presence of BOLD signal flow between $\tau = -14$ and $\tau = -7$. Furthermore, we can observe a deactivating wave starting in the peripheral section of early visual cortex (which, in our definition, is part of the TPN), spreads towards the fovea (which, in our definition, is part of the TNN), and is followed by a similar activating wave. These waves were observed in all four scanning sessions (see SM). These observations suggest that, rather than acting independently, TNN and TPN are part of a more comprehensive dynamical organization in time and space that links TNN and TPN through traveling waves within early visual cortex. The hypothesis that functional networks are coordinated through BOLD signal flow has been proposed before and explored using different analysis methods (Majeed et al., 2009, 2011; Mitra et al., 2014; Matsui et al., 2016).

3.5. Traveling BOLD waves in early visual cortex

In this section we will have a closer look at the traveling waves in early visual cortex observed in Section 3.4. We thus restrict the flow map of TFM 5 to early visual cortex i.e. V1–V4 according to (Glasser et al., 2016a). These regions-of-interest (ROI's) are shown in Fig. 5A on cortical flat maps. Each of the four ROI's is organized retinotopically. Specifically, in each hemisphere, each of the four ROI's corresponds to an elongated strip of cortex, where the upper and lower half of the strip correspond to the lower and upper half of the visual field, respectively. The center of the visual field (the fovea) corresponds to the point where the strips meet (at about the heights of the indentation in the flat maps) and eccentricity increases along the strips (orientation within the visual field is organized perpendicularly to the strips).

Fig. 5B shows the restricted flow map of TFM 5, taken from scanning session RL₁, for lags ranging between ± 14 samples. The flow maps extracted from the other three scanning sessions are similar and are displayed in Supplementary Figs. 9–11. The figure shows a deactivating wave that starts in the periphery of the visual field, propagates towards the fovea, and is followed by a similar activating wave. Comparing the first and last frame, the map is seen to correspond to a bit less than one oscillation cycle. To obtain a rough estimate of the propagation speed, we assume that one oscillation cycle corresponds to 36 samples, which corresponds to a frequency of $f \approx 0.04$ Hz ($TR = 0.72$ seconds). Furthermore, Fig. 5B shows that half the wavelength corresponds to the extent of V1, taken along the eccentricity-axis, which is about 60 mm. Thus, the wavelength is about $\lambda \approx 120$ mm. The propagation speed therefore is about $v = \lambda f \approx 0.05$ m/s, which lies well within the range of haemodynamic traveling waves in V1 elicited by visual stimuli (Aquino et al., 2012; Pang et al., 2018), which propagate with speeds in the range 0.02–0.12 m/s but only over very short distances (5–10 mm). Given this correspondence in propagation speeds, we speculate that the physiological mechanisms underlying induced and spontaneous BOLD signal propagation are similar and can be modeled by the spatiotemporal hemodynamic response function formulated in (Aquino et al., 2012; Pang et al., 2018). It would be interesting to assess how the spatiotemporal deconvolution method proposed in (Pang et al., 2018) can be applied to the extracted propagation maps to reconstruct the underlying neural traveling wave.

The flow map shows that the BOLD signal waves are synchronized across hemispheres and across the upper and lower visual field, which implies that the wave propagates from the periphery towards the center

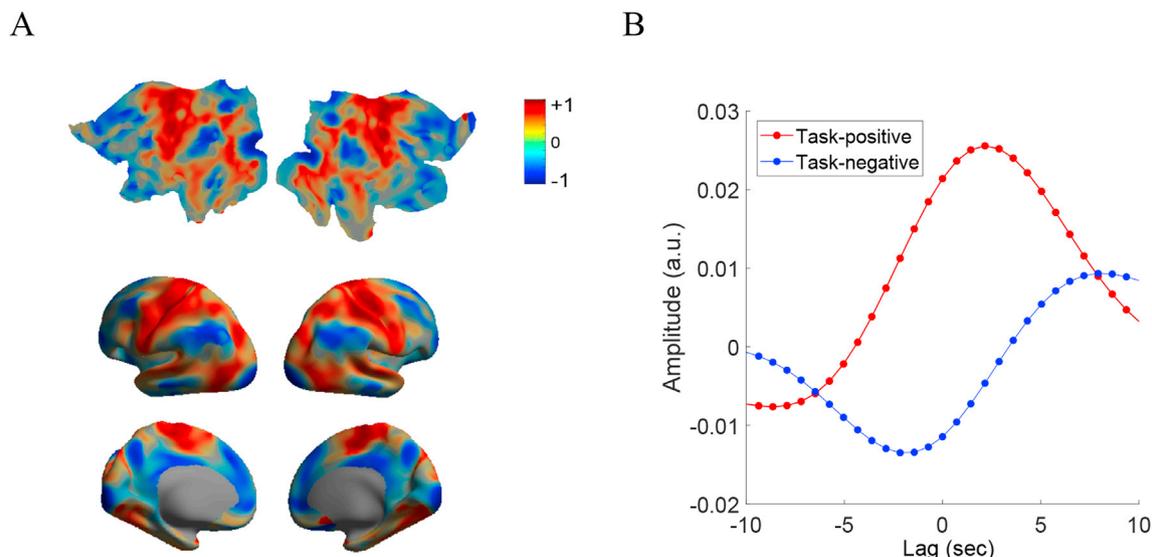


Fig. 4. Coordination between task-positive and task-negative resting-state networks. A. Phase-difference between the second and the third harmonic of the latency signals of TFM 5 as a function of location. The phase-differences have been scaled to lie between -1 and 1 . B. Latency signals of TFM 5, averaged over all cortical vertices in the task-positive network (TPN) (red) and the task-negative network (TNN) (blue).

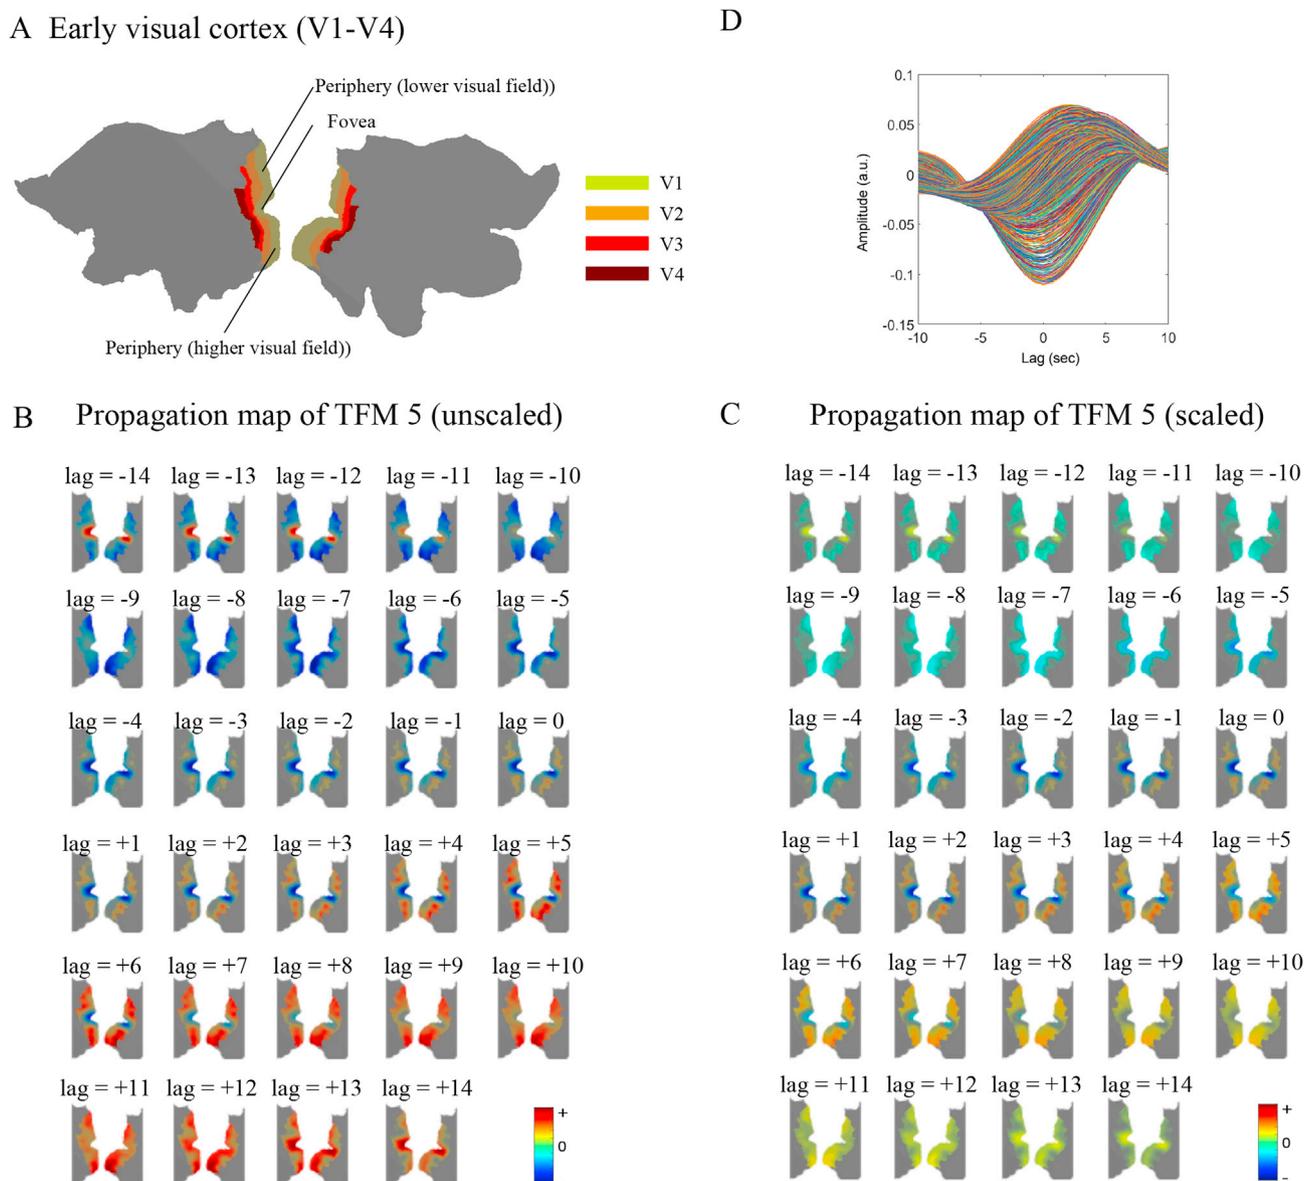


Fig. 5. Traveling BOLD signal waves in early visual cortex. A. Cortical flat map indicating visual areas V1-V4 according to (Glasser et al., 2013). B. Flow map of TFM 5 restricted to V1-V4. Lags range from $\tau = -14$ samples to $\tau = +14$ samples, relative to the mode's peak-activation at $\tau = 0$. Colormaps are unscaled so that amplitudes cannot be compared across frames. Time is in the reading direction. C. Same as in B. but with scaled colorbars so that amplitudes can directly be compared across frames. D. The 500 highest-variance latency signals in the flow map of TFM 5 restricted to V1-V4.

of the visual field simultaneously in all four quadrants (left, right, up, down). Moreover, the wave seems to be synchronized across V1-V4, although this can only be established with certainty by using retinotopic mapping. In any case, locations within each of the ROI's that correspond to a particular eccentricity, are (de)activated *simultaneously* and locations with different eccentricities are (de)activated *sequentially*. The dispersion of peak-times of the wave across different cortical locations is visible in Fig. 5D, which shows the 500 highest-variance latency signals of the flow map restricted to V1-V4. The figure shows that peak-times can differ between locations by as much as about 5 s. Moreover, since the same waves can be observed in all four scanning sessions, which have different phase-encoding directions, they cannot be due to the omission of slice-timing correction, since the latter can only lead to differences in peak-times < 1 sample. The synchronized waves explain the recent observation that iso-eccentric locations within early visual regions are functionally connected with zero-lag (Arcaro et al., 2015), because such regions are deactivated at the *same* phase of the wave. It also explains why no zero-lag functional correlations are observed between

locations with different eccentricities (Arcaro et al., 2015; Dawson et al., 2016), because such regions are deactivated at *different* phases of the wave. Our observation provides more evidence for the hypothesis that zero-lag functional connectivity arises from signal flow (Mitra et al., 2014; Matsui et al., 2016).

3.6. Lagged correlations between temporal functional modes

In Section 2.6 we showed that the presence of flow about a TFM implies lagged correlations between the mode and some of the other modes. In this section we analyze the lagged correlation structure of the extracted TFM's and relate it to the observed flow about Mode 5. Because any of the nine extracted reliable TFM's can be correlated to TFM's that were not considered reliable, analyzing the correlation structure requires considering all 20 initially extracted TFM's (see Section 2.3). For the analysis to be valid, all TFM's need to be uncorrelated at lag zero (see Section 2.6) which was the case (all zero lag correlations from all sessions are on the order of 10^{-15}). Fig. 6A (left panel) shows the lag-averaged

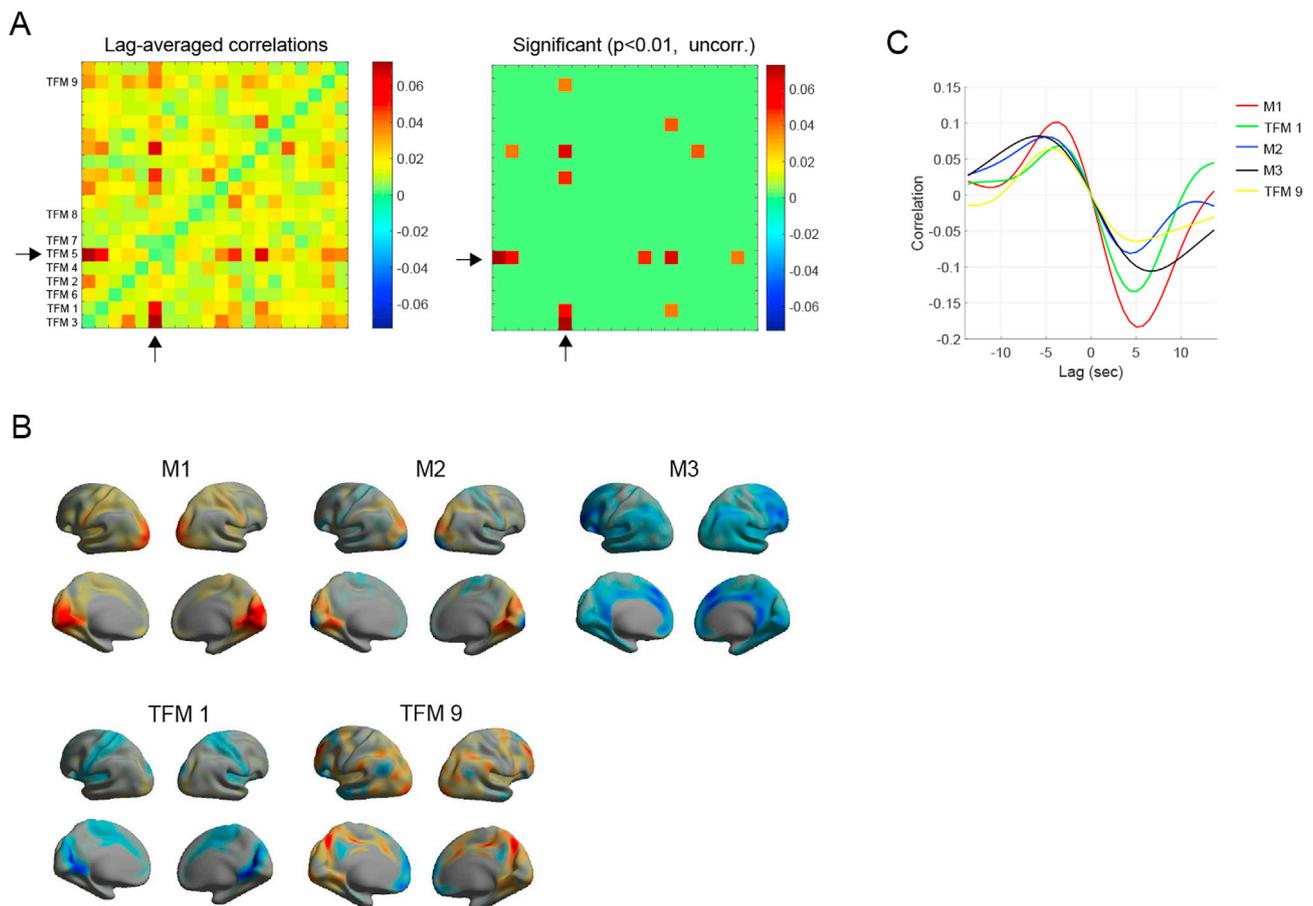


Fig. 6. Lagged correlations between temporal functional modes. **A.** Left panel: Lag-averaged correlations between all 20 TFM's extracted from scanning session RL_1 . Right panel: Same as **A.** but showing only the significant correlations ($p < 0.01$, uncorrected). Black arrow designate TFM 5. **B.** The temporal functional modes that are significantly lag-correlated with TFM 5. Two of them correspond to reliable TFM's (reliable TFM 1 and TFM 9). The other three are labeled M_1 , M_2 , and M_3 . **C.** Cross-correlation functions between TFM 5 and the five significantly lag-correlated TFM's.

correlation matrix of the 20 TFM's extracted from scanning session RL_1 . The nine reliable TFM's are designated on the vertical axis. So, for example, reliable TFM 9 corresponds to the 19-th TFM of scanning session RL_1 .

To assess significance, we generated 99 coherent phase-randomized surrogate copies of the data matrix, applied temporal ICA to extract 20 TFM's, and calculated their lag-averaged correlation matrix. To avoid having to pair the 20 TFM's with those extracted from the unrandomized data, we selected the maximal value from each surrogate correlation matrix. If the observed lag-averaged correlation between a given pair of TFM's exceeded all 99 surrogate values, it was considered significant ($p < 0.01$, uncorrected). Because correction for multiple tests (one test for every ordered pair of TFM's) requires the generation of an unfeasible number of surrogate data matrices, we gained additional confidence in the results by demonstrating reproducibility over scanning sessions. Fig. 6A (right panel) shows the significant lag-integrated correlations between the 20 modes. It demonstrates that several modes are significantly correlated and hence that TFM's are not temporally independent when non-zero lags are taken into account. Supplementary Figs. 12–14 demonstrate that these findings are fairly robust across scanning sessions.

We focus on correlations with reliable TFM 5, which corresponds to the 6th TFM of scanning session RL_1 . This TFM is designated by black arrows in Fig. 6A. The figure shows that TFM 5 is significantly correlated with 5 other TFM's. Fig. 6B shows these TFM's. Two of these TFM's are in fact reliable and correspond to reliable TFM 1 and TFM 9. We will refer to the three non-reliable TFM's as M_1 , M_2 , and M_3 . M_1 shows strong activation in early visual cortex and weaker activations in auditory,

somatosensory and motor cortex. Incidentally, M_1 is reliable in the sense of having a spatial correlation > 0.7 with modes in the other three scanning sessions, but was not regarded as reliable by our criterion (see Section 2.3) because it had a spatial correlation > 0.7 with more than one other mode. M_2 shows activation in the peripheral section of early visual cortex and deactivation in the foveal section. It resembles TFM 5 but does not involve other sensory modalities and motor cortex and shows a sharper contrast (activation-deactivation) between the foveal and peripheral sections of early visual cortex. M_3 shows strong activation in cingulate and dorsolateral prefrontal cortices and weaker activation in inferior parietal and insular cortices and thus largely overlaps with the DMN.

Fig. 6C shows the cross-correlation functions between TFM 5 and the five significantly correlated TFM's. Lags are relative to the peak-activation of TFM 5. Note that the significant lag-integrated correlations in Fig. 6A (right panel) were obtained by averaging the respective (absolute) cross-correlation functions over lags (see Section 2.6). To simplify the analysis, the signs of the cross-correlation functions are chosen such that the modes activate before and deactivate after lag zero. We make several observations. First, the cross-correlation functions vanish at lag zero, which reflects the absence of zero-lag correlations with TFM 5. Second, their peak-activation times, relative to the peak-activation time of TFM 5 differ from zero (up to about ± 5 seconds) and also from each other, which provides evidence that, relative to the peak-activation of TFM 5, the modes activate in a stereotypical temporal progression. We do remark, however, that the omission of slice-timing correction (see Section 2.1) might be responsible for some of these latencies, namely those ≤ 1 samples.

3.7. A simple model system

In this study we have provided evidence for entanglement of ongoing BOLD fluctuations in time and space (i.e. flow) in the temporal vicinity of TFM peak-activations and demonstrated significant lagged-correlations between TFM's. Furthermore, in Section 2.6 we have used a simple mathematical model to argue that the existence of flow and lagged-correlations within the data are equivalent. Both arise from the latency structure inherent in the data. In this section, we use a simple model system to make explicit the relationships between latency structure, flow, and lagged correlations, which will provide intuition about the findings reported in this study and about TFM's in general.

We consider an oscillatory traveling wave $X(p, t)$ on the line segment $[0, L]$, interpreted as a one-dimensional “cortex” of length L . The wave has temporal frequency f , wavelength λ , and travels from left to right:

$$X(p, t) = \cos(2\pi(ft - p/\lambda) + \phi), \quad (20)$$

where ϕ is a phase-offset and where $p \in [0, L]$ and t denote position and time, respectively. Fig. 7A shows X as a function of p and t , displaying oscillatory traveling waves that propagate from left to right. Note that the dynamics are entangled in space and time and are therefore inherently spatiotemporal.

To obtain the TFM's, we express $X(p, t)$ in a basis of temporally independent modes. Since the modes will be oscillatory as well and have the same temporal frequency f , the TFM's can only be independent if their phase-difference equals $\pm\pi/2$. Concretely, $X(p, t)$ can be written as a superposition of two TFM's:

$$X(p, t) = a_1(t)f_1(p) + a_2(t)f_2(p), \quad (21)$$

where $f_1(p) = \cos(2\pi p/\lambda + \phi)$ and $f_2(p) = \sin(2\pi p/\lambda + \phi)$ are the TFM's at location p and $a_1(t) = \cos(\omega t)$ and $a_2(t) = \sin(\omega t)$ are their respective time-courses. Note that f_1 and f_2 are indeed independent:

$$\mathbb{E}[a_1(t)a_2(t)] = \frac{1}{T} \int_0^T a_1(t)a_2(t)dt = \frac{1}{2T} \int_0^T \sin(2\omega t)dt = 0, \quad (22)$$

where $T = 2\pi/\omega$ is the angular frequency of the oscillations. The last equality holds because the period of $\sin(2\omega t)$ equals $T/2$. The spatial profiles and time-courses of the TFM's are shown in Fig. 7B and C, respectively. Note that, similar to experimental TFM's (Smith et al., 2012), one of the model TFM's has “activated” as well as “deactivated”

regions and overlap in space, and that the latency between their time-courses is such that zero-lag correlations vanish.

In Supplementary Text I we derive that the modes' flow maps Λ_1 and Λ_2 are given by

$$\Lambda_1(p, \tau) = \frac{1}{2}X(p, \tau), \quad (23)$$

and

$$\Lambda_2(p, \tau) = \frac{1}{2}X(p, \tau + T/4). \quad (24)$$

The modes are hence proportional to the traveling wave X about two distinct time-points (0 and $T/4$, respectively). When applied to this model system, temporal ICA thus breaks down the traveling wave into two standing waves, both of which are “snapshots” of the traveling wave at distinct time-points. In Supplementary Text I we also derive that the cross-correlation function $\gamma_{12}(\tau)$ between the mode's time-courses $a_1(t)$ and $a_2(t)$ is given by

$$\gamma_{12}(\tau) = -\sin(\omega\tau), \quad (25)$$

which shows that zero-lag correlations vanish and that the modes are correlated at non-zero lags.

4. Discussion

Summary and main findings. Whereas most resting-state BOLD-fMRI studies currently do not exploit the latency structure within the data, the existence of latencies in resting-state BOLD-fMRI data is receiving increased attention due to faster acquisition rates and the development of new analysis methods, fueled by the discovery of propagating activity patterns and latency structure that are consistent across scans, participants, and data-sets (Majeed et al., 2009, 2011; Mitra et al., 2014, 2015a; Siegel et al., 2016; Amemiya et al., 2016; Matsui et al., 2016; Raatikainen et al., 2017; Gravel et al., 2017). Furthermore, there is both indirect as well as direct evidence that the latency structure of resting-state BOLD fluctuations can only partly be explained by spatial variability in haemodynamic responses and are of electrophysiological origin (Mitra et al., 2014; Thompson et al., 2014; Amemiya et al., 2016; Matsui et al., 2016; Vanni et al., 2017). It is, however, still unclear how BOLD propagation is related to the more classical notion of resting-state networks (RSN's), based on spatial ICA or seed-based FC, although several studies have

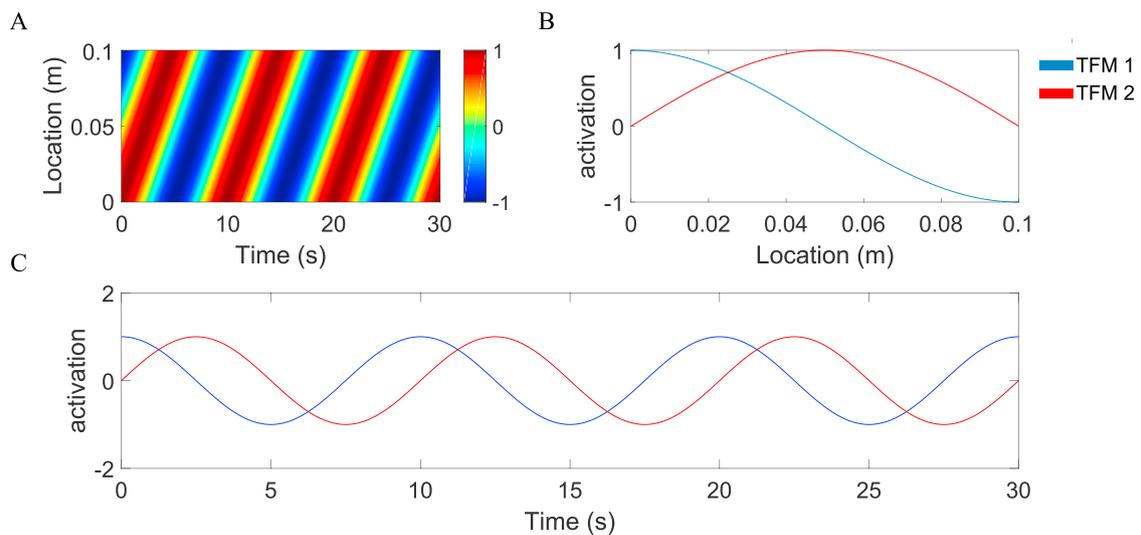


Fig. 7. A simple model system. A. Simulated oscillatory activity on a one-dimensional cortex of length $L = 10$ cm as a function of time and location. The oscillations propagate from left to right. Temporal frequency and wavelength were set to $f = 0.1$ Hz and $\lambda = 20$ cm, respectively. B. Temporal functional modes (TFM's). C. Time-courses of the TFM's in panel B.

started to explore their relationship (Mitra et al., 2015a; Matsui et al., 2016; Vanni et al., 2017). Our study contributes to this line of research by presenting a method for characterizing the spatiotemporal dynamics of ongoing BOLD signals about a given functional network and by demonstrating its merit using experimental BOLD-fMRI data. By applying the method to functional networks obtained through temporal ICA i.e. to temporal functional modes (TFM's) (Smith et al., 2012), extracted from resting-state BOLD-fMRI data from 94 participants provided by the Human Connectome Project (Glasser et al., 2013), several TFM's were found to be embedded in BOLD signal flow and detailed analysis of the BOLD fluctuations about one of the modes revealed a linkage between task-positive and task-negative regions by traveling BOLD waves in early visual cortex that follow eccentricity gradients. Our study is significant in that it provides further evidence that functionally distinct networks are part of a larger dynamical organization, constituted by the brain's latency structure (Majeed et al., 2009; Mitra et al., 2014, 2015a, 2018; Thompson et al., 2014; Matsui et al., 2016; Vanni et al., 2017).

Interpretation of TFM amplitudes. We have argued for the usefulness of the proposed method by applying it to a particular kind of cortical states, namely to temporal functional modes (TFM's) (Smith et al., 2012). The first result of this application is a deeper understanding of TFM amplitudes and how they can be interpreted. The common view is that TFM amplitudes measure the relative contributions of different voxels and that negative amplitudes reflect functional incompatibility or suppression (Smith et al., 2012). This view presupposes that the peak-activations of a given TFM coincide with simultaneous (de)activation at the contributing voxels. Although this is the case for most of the TFM's, we have demonstrated that this need not always be true. Specifically, this assumption is correct only in the absence of flow about the TFM. If the TFM is embedded in a flow, ongoing BOLD fluctuations can peak before as well as after the TFM's peak-activation and, consequently, the identity of amplitude and contribution strength is broken as timing starts to play a role. Thus, besides peak-deactivation, a negative amplitude at a particular voxel can reflect earlier activation and later deactivation or *vice versa* and which interpretation is correct depends on the sign of the BOLD signal's slope i.e. whether the signal is rising or falling. The point is that the correct conclusion can only be drawn by analyzing the dynamics of the BOLD signal in an entire temporal window and not by observing a single data frame. These considerations apply not only to TFM's, but to cortical states in general, such as RSN's, (innovative) co-activation patterns (Liu and Duyn, 2013; Liu et al., 2013; Karahanoglu and Van De Ville, 2015) among others (Yeo et al., 2014). Naturally, BOLD signal flow can only be detected if acquisition rates are high enough or if suitable interpolation is used (Mitra et al., 2014).

Anti-correlations between task-positive and task-negative regions. Analysis of the most strongly flowing TFM showed near anti-correlations between task-negative (i.e. the DMN) and task-positive regions, in line with studies using seed-based correlation analysis (Fransson, 2005; Fox et al., 2005) and with studies that exploit the latency structure within the data (Majeed et al., 2011; Vidaurre et al., 2017). Interestingly, we observed a delay of about 3 s between the deactivation of the TNN and activation of the TPN (see Fig. 4B), which vanished when the global signal was subtracted from the data, whereas the visual traveling waves could still be observed (results not shown). What has not been observed before, however, is that ongoing BOLD fluctuations in task-positive and task-negative regions are integrated through traveling waves in early visual cortex. Specifically, in the delay period between deactivation of TNN and activation of TPN, a traveling wave in early visual cortex (V1-V4) could be observed that starts from the periphery of the visual field and propagates towards the fovea. Although we cannot rule out that fixation on a cross-hair contributed to these observations, it is interesting to note that delayed interactions between BOLD signals in early visual cortex have been observed in resting-state protocols in which the participants have their eyes closed (Gravel et al., 2017). A possible functional interpretation is that periphery-to-fovea waves reflect task related changes in the power of alpha oscillations as they have been implicated

in surround suppression and facilitation in the fovea and are associated with concurrent deactivations in BOLD activity (Harvey et al., 2013). In (Majeed et al., 2011), propagation of BOLD signals related to ongoing BOLD fluctuations in TNN and TPN was observed in several cortical regions, but visual traveling waves were not reported. Perhaps the absence of visual waves is due to the lower accuracy of volumetric preprocessing (including registration to MNI space) of the data in (Majeed et al., 2011). However, in our data we did not observe propagation within other cortical regions. Although such waves might have averaged out over subjects, as higher-order functional areas display more inter-subject variability, in single-subject data, we also did not observe waves in other cortical regions. We do not rule out, however, that the HCP data used in our study might contain additional, more subtle or transient sources of BOLD signal propagation in other cortical regions as observed in (Majeed et al., 2011) that are undetected by our method. In any case (Majeed et al., 2011), and our study provide evidence that ongoing BOLD fluctuations in TNN and TPN comprise an intrinsically spatiotemporal process and, as such, cannot be disentangled into separate spatial and temporal components.

Lagged coordination between temporal functional modes. We have shown mathematically that the existence of flow between (some of the) TFM's is equivalent to the existence of lagged correlations between (some of the) TFM's and have confirmed this in experimental data using lag-averaged correlations and appropriate surrogate data. In particular, TFM 5 was found to be significantly lag-correlated with five other modes, which display a temporal progression in peak-activations, relative to the peak-activation of TFM 5. Lagged-correlation analysis thus enables conceptualizing the flow about functional modes as arising from precisely timed lagged-interactions between the mode and other modes. The existence of flows and, equivalently, lag-averaged correlations, implies a non-separable latency structure within the data i.e. entanglement of ongoing BOLD fluctuations in time and space. We have used a simple mathematical model of spatiotemporal activity that reproduces and clarifies the observed relationships between latency structure, flow, and lagged correlations. Although latency structure in resting-state BOLD-fMRI has been studied before (Majeed et al., 2009; Mitra et al., 2014, 2015a; Thompson et al., 2014; Matsui et al., 2016), an important contribution of our study is to exploit this structure to better understand how different aspects of resting-state BOLD-fMRI dynamics are reflected in the spatial and temporal features of TFM's (Smith et al., 2012).

Statistical pitfalls in latency analysis. In the absence of a baseline or experimental contrast, which is frequently the case in resting-state studies, significance of an observed effect is usually tested by constructing randomized data to approximate the null-distribution of the used test-statistic. A potential pitfall is the use of inappropriate null-data i.e. null-data that corresponds to a different null-hypothesis than intended. An educative example comes from the field of dynamic FC, in which non-stationarity of resting-state functional connectivity is investigated. In this field, the use of inappropriately constructed null-data is widespread and has led to incorrect claims (Hindriks et al., 2016). A frequently encountered misuse is to randomize the data such that *all* FC is destroyed and not only dynamic FC. This randomization corresponds to the (unintended) null-hypothesis that stationary FC is absent. Rejecting the null-hypothesis, therefore, does not allow to conclude that FC is dynamic, but merely that FC is non-zero. More examples can be found in (Hindriks et al., 2016). The use of inappropriate null-data is encountered in studies on non-stationarity of resting-state brain dynamics in general, including micro-state analysis, co-activation patterns (CAP's) (Liu and Duyn, 2013), and quasi-periodic patterns (QPP's) (Majeed et al., 2011). In (Liu and Duyn, 2013), non-stationarity of resting-state BOLD-fMRI is tested for by generating null-data that lacks both auto- and cross-correlations. In (Majeed et al., 2011), incoherent phase-randomization (IPR) is used to argue that resting-state BOLD-fMRI contain reoccurring spatiotemporal patterns (QPP's) that are linked to fluctuations in FC. This is inappropriate because IPR destroys *all* cross-correlations within the data (Schreiber and Schmitz, 2000). This

does not mean that the data does not contain QPP's, as they clearly do (Majeed et al., 2009, 2011; Keilholz et al., 2016), have shown to be reproducible (Majeed et al., 2009), and correlate with local field potentials (Thompson et al., 2014). It also does not mean that QPP's have no functional role or are otherwise uninteresting. It merely means that, until rejected by using appropriate null-data, which in this context is coherent phase-randomization (Hindriks et al., 2016), QPP's are consistent with the data having static FC.

To test for the presence of flow in resting-state data, null-data should ideally have the same auto- and cross-covariance functions as the experimental data, except that latencies have been destroyed i.e. the cross-covariance functions have been shifted in time to peak at lag zero. We have approximated such ideal surrogate data by generating copies of the experimental data that have symmetric cross-covariance functions and approximately the same auto-correlation functions. This was done by taking the real part of the Fourier transformed data, adding random but coherent phases, and transforming back to the time domain. This data is not ideal, however, as cross-covariance functions can become deformed instead of only translated. Furthermore, asymmetry in cross-covariance functions gets destroyed, which implies that rejecting the null-hypothesis can also reflect that the data is time-irreversible i.e. have asymmetric covariance structure, which is a weaker property than exhibiting latencies. Fortunately, cross-covariance functions of resting-state BOLD-fMRI data are rather symmetric, at least for not too large lags. In relation to this, in (Mitra et al., 2014) latencies between BOLD signals are estimated by parabolic interpolation of the estimated cross-covariance function at lags -1 , 0 , and $+1$, so that the estimated lags effectively measure the asymmetry of the covariance function. It is only for symmetric covariance functions that asymmetry of the estimated covariance function implicates a non-zero lag. Another issue is at which step in the processing pipeline the randomization should be done. Generally, randomizing at later stages yields more spurious results since this tends to underestimate the sampling variance of the used test-statistic. This is illustrated by the following. In testing for significant lagged-correlations between TFM time-courses, which were extracted using ICA, we randomized the pre-processed BOLD-fMRI time-series and subsequently extracted the surrogate TFM's. Alternatively, testing by directly using the time-courses of the extracted TFM's, thus after the application of ICA, yielded many more significant correlations (results not shown). Randomization of the data in a later stage of processing or analysis typically yields more significant but nevertheless spurious effects and should therefore be done with caution.

Slice-timing correction. In this study we used the FIX-denoised-compact data from the Human Connectome Project (HCP) (Smith et al., 2013), the pre-processing of which did not include slice-timing correction. A consequence is that observed latencies between BOLD signals equal to 1 TR (720ms in our data) might result from the acquisition protocol and, as such, have no physiological significance. This applies even if the latencies are statistically significant, because the null-distribution is approximated by randomizing the same data. Another consequence is that the accuracy of observed latencies is ± 1 TR. For example, if a latency of 3 TR's is observed, one can only conclude that the underlying physiological latency lies between 2 and 4 TR's (ignoring sampling variability in the used estimator). The latencies observed in our study, however, range up to ± 5 TR's and can therefore not entirely be explained by differences in slice-timing, although the propagations maps might be a bit distorted. The (possible) distortion can be thought of as due to inaccuracies in the relative timings of the TFM's (see Section 3.6 and Fig. 6C) which are about 1–2 TR's in absolute value. The main features of the propagation map analyzed in Sections 3.4 and 3.5 i.e. visual traveling waves and their relationship with task-positive and task-negative networks, however, are due to the timings of the TFM's relative to the reference TFM, which range up to 7 TR's in absolute value and can therefore not be due to differences in slice-timing.

Time-resolved latency analysis. In the present study we proposed a method for characterizing the flow of ongoing BOLD signals about a given brain state. The basic mathematical tool in our analysis is the flow kernel, which is defined in terms of cross-covariance functions that are estimated from the entire observation period. This presupposes that ongoing BOLD fluctuations are stationary and our method, as well as several other methods (Mitra et al., 2014, 2015a), therefore is only capable of extracting the most dominant flow pattern. Recent modeling work, however, using large-scale dynamical models on the human connectome, suggests that cortical flow patterns are multistable and exhibit phase-transitions (Roberts et al., 2019). To study such phenomena, a dynamic analysis in which no assumptions about stationarity are made, is required (Majeed et al., 2009, 2011; Hindriks et al., 2014; Thompson et al., 2014; Matsui et al., 2016). One way of extending a stationary latency analyses is to estimate the cross-covariance functions using a sliding-window. Apart from statistical difficulties with this approach (Hindriks et al., 2016), at any given point in time, typically a linear superposition of wave patterns is observed, and it is unclear how the contributing patterns can be retrieved. The methods proposed in (Majeed et al., 2009, 2011; Hindriks et al., 2014; Thompson et al., 2014; Matsui et al., 2016) also suffer from this drawback. In our study, separate propagation patterns were reconstructed rather *ad hoc*, namely by first applying temporal ICA and subsequently reconstructing the flow about the extracted states. Moreover, this approach might not be feasible for use in combination with sliding-windows because of the relatively large number of samples that is needed. Further research is needed to develop time-resolved analysis methods that are able to reconstruct transient dynamical phenomena from linearly mixed patterns.

Acknowledgments

The authors have no financial or non-financial competing interests. R.H. was funded by the European Research Council (Advanced Grant DYSTRUCTURE No.295129) and by the NWO-Wiskundeclusters grant nr. 613.009.105. G.D. was funded by the Spanish Research Project PSI2016-75688-P (AEI/FEDER, EU), by the European Unions Horizon 2020 Research and Innovation Programme under grant agreements n. 720270 (HBP SGA1) and n. 785907 (HBP SGA2), and by the Catalan AGAUR Programme 2017 SGR 1545. Nicolas Gravel was supported by the (Chilean) National Commission for Scientific and Technological Research (BECAS CHILE) and the Graduate School for Medical Sciences (GSMS) of the University Medical Center Groningen (UMCG).

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.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.neuroimage.2019.06.007>.

References

- Amemiya, S., Takao, H., Hanaoka, S., Ohtomo, K., 2016. Global and structured waves of rs-fMRI signal identified as putative propagation of spontaneous neural activity. *NeuroImage* 133, 331–340.
- Aquino, K.M., Schira, M.M., Robinson, P.A., Drysdale, P.M., Breakspear, M., 2012. Hemodynamic traveling waves in human visual cortex. *PLoS Comput. Biol.* 8 (3).
- Arcaro, M.J., Honey, C.J., Mruzczek, R.E.B., Kastner, S., Hasson, U., 2015. Widespread correlation patterns of fMRI signal across visual cortex reflect eccentricity organization. *Elife* (4), 1–28, 2015.
- Beckmann, C.F., DeLuca, M., Devlin, J.T., Smith, S.M., 1457. Investigations into resting-state connectivity using independent component analysis. *Philos. Trans. R. Soc. B Biol. Sci.* 360, 1001–1013, 2005.
- Benucci, A., Frazor, R.A., Carandini, M., Jul 2007. Standing waves and traveling waves distinguish two circuits in visual cortex. *Neuron* 55 (1), 103–117.

- Biswal, B.B., Mennes, M., Zuo, X.-N., Gohel, S., Kelly, C., Smith, S.M., Beckmann, C.F., Adelstein, J.S., Buckner, R.L., Colcombe, S., Dogonowski, A.-M., Ernst, M., Fair, D., Hampson, M., Hoptman, M.J., Hyde, J.S., Kiviniemi, V.J., Kötter, R., Li, S.-J., Lin, C.-P., Lowe, M.J., Mackay, C., Madden, D.J., Madsen, K.H., Margulies, D.S., Mayberg, H.S., McMahon, K., Monk, C.S., Mostofsky, S.H., Nagel, B.J., Pekar, J.J., Peltier, S.J., Petersen, S.E., Riedel, V., Rombouts, S.A.R.B., Rypma, B., Schlaggar, B.L., Schmidt, S., Seidler, R.D., Siegle, G.J., Sorg, C., Teng, G.-J., Veijola, J., Villringer, A., Walter, M., Wang, L., Weng, X.-C., Whitfield-Gabrieli, S., Williamson, P., Windischberger, C., Zang, Y.-F., Zhang, H.-Y., Castellanos, F.X., Milham, M.P., 2010. Toward discovery science of human brain function. *Proc. Natl. Acad. Sci.* 107 (10), 4734–4739.
- Cecchi, G.A., Rao, A.R., Centeno, M.V., Baliki, M., Apkarian, A.V., Chialvo, D.R., 2007. Identifying directed links in large scale functional networks: application to brain fMRI. *BMC Cell Biol.* 8 (Suppl 1(Suppl 1)):S5.
- Contreras, D., Llinas, R., 2001. Voltage-sensitive dye imaging of neocortical spatiotemporal dynamics to afferent activation frequency. *J. Neurosci.* 21 (23), 9403–9413.
- 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. *Proc. Natl. Acad. Sci.* 103 (37), 13848–13853.
- Dawson, D.A., Lam, J., Lewis, L.B., Carbonell, F., Mendola, J.D., Shmuel, A., 2016. Partial correlation-based retinotopically organized resting-state functional connectivity within and between areas of the visual cortex reflects more than cortical distance. *Brain Connect.* 6 (1), 57–75.
- 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. *Proc. Natl. Acad. Sci. U.S.A.* 102 (27), 9673–9678.
- 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 (6), 3270–3283.
- Fransson, P., 2005. Spontaneous low-frequency BOLD signal fluctuations: an fMRI investigation of the resting-state default mode of brain function hypothesis. *Hum. Brain Mapp.* 26 (1), 15–29.
- Garg, R., Cecchi, G.A., Rao, A.R., 2011. Full-brain auto-regressive modeling (FARM) using fMRI. *Neuroimage* 58 (2), 416–441.
- 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 (7615), 171–178.
- Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Van Essen, D.C., Jenkinson, M., 2013. The minimal preprocessing pipelines for the Human Connectome Project. *Neuroimage* 80, 105–124.
- Glasser, M.F., Smith, S.M., Marcus, D.S., Andersson, J., Auerbach, E.J., Behrens, T.E.J., Coalson, T.S., Harms, M.P., Jenkinson, M., Moeller, S., Robinson, E.C., Sotiropoulos, S.N., Xu, J., Yacoub, E., Ugurbil, K., Van Essen, D.C., 2016. The Human Connectome Project's neuroimaging approach. *Neurosci.* (9), 1175–1187 (in press).
- Golomb, D., Amitai, Y., 1997. Propagating Neuronal Discharges in Neocortical Slices : Computational and Experimental Study. *Am. Physiol. Soc.*, pp. 1199–1211
- Gravel, N., Renken, R.J., Harvey, B.M., Deco, G., Cornelissen, F.W., Gilson, M., 2017. Propagation of BOLD Activity Reveals Directed Interactions across Human Visual Cortex.
- Griffa, A., Ricaud, B., Benzi, K., Bresson, X., Daducci, A., Vanderheyne, P., Thiran, J.P., Hagmann, P., 2017. Transient networks of spatio-temporal connectivity map communication pathways in brain functional systems. *Neuroimage* 155, 490–502.
- Griffanti, Ludovica, Salimi-Khorshidi, Gholamreza, Beckmann, Christian F., Auerbach, Edward J., Douaud, Gwenaelle, Sexton, Claire E., Zsoldos, Eniko, Ebmeier, Klaus P., Filippini, Nicola, Mackay, Clare E., Moeller, Steen, Xu, Junqian, Yacoub, Essa, Baselli, Giuseppe, Ugurbil, Kamil, Miller, Karla L., Smith, Stephen M., 2014. ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging. *Neuroimage* 95, 232–247.
- Harvey, B.M., Vansteensel, M.J., Ferrier, C.H., Petridou, N., Zuiderbaan, W., Aarnoutse, E.J., Bleichner, M.G., Dijkerman, H.C., Van Zandvoort, M.J.E., Leijten, F.S.S., Ramsey, N.F., Dumoulin, S.O., 2013. NeuroImage Frequency specific spatial interactions in human electrocorticography. V1 alpha oscillations reflect surround suppression 65, 424–432.
- Hindriks, R., Adhikari, M.H., Murayama, Y., Ganzetti, M., Mantini, D., Logothetis, N.K., Deco, G., 2016. Can sliding-window correlations reveal dynamic functional connectivity in resting-state fMRI? *Neuroimage* 127, 242–256.
- Hindriks, R., van Putten, M.J.A.M., Deco, G., dec 2014. Intra-cortical propagation of EEG alpha oscillations. *Neuroimage* 103, 444–453.
- Hyvärinen, A., 1999. Survey on independent component analysis. *Neural Comput. Surv.* 10 (3), 626–634.
- Karahanoglu, F.I., Van De Ville, D., 2015. Transient brain activity disentangles fMRI resting-state dynamics in terms of spatially and temporally overlapping networks. *Nat. Commun.* 6, 7751.
- Keilholz, S.D., Billings, J.C.W., Wang, K., Abbas, A., Hafener, C., Pan, W.-J., Shakil, S., Nezafati, M., 2016. Multiscale network activity in resting state fMRI. 2016 38th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. 61–64.
- Larson-Prior, L.J., Oostenveld, R., Della Penna, S., Michalareas, G., Prior, F., Babajani-Feremi, A., Schoffelen, J.M., Marzetti, L., de Pasquale, F., Di Pompeo, F., Stout, J., Woolrich, M., Luo, Q., Bucholz, R., Fries, P., Pizzella, V., Romani, G.L., Corbetta, M., Snyder, A.Z., 2013. Adding dynamics to the human connectome project with MEG. *Neuroimage* 80, 190–201.
- Liu, T.T., Nalci, A., Falahpour, M., 2017. The global signal in fMRI: nuisance or Information? *Neuroimage* 150, 213–229.
- Liu, X., Chang, C., Duyn, J.H., 2013. Decomposition of spontaneous brain activity into distinct fMRI co-activation patterns. *Front. Syst. Neurosci.* 7 (December), 101.
- Liu, X., Duyn, J.H., mar 2013. Time-varying functional network information extracted from brief instances of spontaneous brain activity. *Proc. Natl. Acad. Sci. U.S.A.* 110 (11), 4392–4397.
- Lubenov, E.V., Siapas, A.G., may 2009. Hippocampal theta oscillations are travelling waves. *Nature* 459 (7246), 534–539.
- Ma, Z., Zhang, N., 2017. Temporal transitions of spontaneous brain activity in awake rats, pp. 1–31.
- Majeed, W., Magnuson, M., Hasenkamp, W., Schwarb, H., Schumacher, E.H., Barsalou, L., Keilholz, S.D., jan 2011. Spatiotemporal dynamics of low frequency BOLD fluctuations in rats and humans. *Neuroimage* 54 (2), 1140–1150.
- Majeed, W., Magnuson, M., Keilholz, S.D., 2009. Spatiotemporal dynamics of low frequency fluctuations in BOLD fMRI of the rat. *J. Magn. Reson. Imaging* 30 (2), 384–393.
- Matsui, T., Murakami, T., Ohki, K., 2016. Transient neuronal coactivations embedded in globally propagating waves underlie resting-state functional connectivity. *Proc. Natl. Acad. Sci.* 113 (23), 6556–6561.
- Mitra, A., Snyder, A.Z., Blazey, T., Marcus, E., 2015a. Lag threads organize the brain's intrinsic activity. *Proc. Natl. Acad. Sci.* 112 (52), E7307–E7307.
- Mitra, A., Snyder, A.Z., Constantino, J.N., Raichle, M.E., Februari 2017. The lag structure of intrinsic activity is focally altered in high functioning adults with autism. *Cereb. Cortex* 27 (2) bhv294.
- Mitra, A., Snyder, A.Z., Hacker, C.D., Raichle, M.E., jun 2014. Lag structure in resting-state fMRI. *J. Neurophysiol.* 111 (11), 2374–2391.
- Mitra, A., Snyder, A.Z., Tagliazucchi, E., Laufs, H., Raichle, M.E., 2015b. Propagated infraslow intrinsic brain activity reorganizes across wake and slow wave sleep. *Elife* 4 (NOVEMBER2015), 1–19.
- Mitra, Anish, Kraft, Andrew, Wright, Patrick, Benjamin, Acland, Snyder, Abraham Z., Rosenthal, Zachary, Czerniewski, Leah, Bauer, Adam, Snyder, Lawrence, Culver, Joseph, Jin, Moo Lee, Marcus, E., Raichle, 2018. Spontaneous infra-slow brain activity has unique spatiotemporal dynamics and laminar structure. *Neuron* 98 (2), 297–305 e6.
- Mohajerani, M.H., Mcvea, D.A., Fingas, M., Murphy, T.H., 2010. Mirrored bilateral slow-wave cortical activity within local circuits revealed by fast bihemispheric voltage-sensitive dye imaging in anesthetized and awake mice. *J. Neurosci.* 30 (10), 3745–3751.
- Pang, J.C., Aquino, K.M., Robinson, P.A., Lacy, T.C., Schira, M.M., 2018. Biophysically based method to deconvolve spatiotemporal neurovascular signals from fMRI data. *J. Neurosci. Methods* 308 (6–20).
- Prichard, D., 1994. Generating surrogate data for time series with several simultaneously measured variables. *Phys. Rev. Lett.* 73 (7), 951–954.
- Raatikainen, V., Huotari, N., Korhonen, V., Rasila, A., Kananen, J., Raitamaa, L., Keinänen, T., Kantola, J., Tervonen, O., Vesa, Kiviniemi, 2017. Combined spatiotemporal ICA (stICA) for continuous and dynamic lag structure analysis of MREG data. *Neuroimage* 148, 352–363. August 2016.
- Roberts, James A., Gollo, Leonardo L., Abeysuriya, Romesh G., Roberts, Gloria, Mitchell, Philip B., Woolrich, Mark W., Breakspear, Michael, 2019. Metastable brain waves. *Nat. Commun.* 10 (1), 1–17.
- Rubino, D., Robbins, K.A., Hatsopoulos, N.G., dec 2006. Propagating waves mediate information transfer in the motor cortex. *Nat. Neurosci.* 9 (12), 1549–1557.
- Salimi-Khorshidi, G., Douaud, G., Smith, S.M., Griffanti, L., Beckmann, C.F., Glasser, M.F., 2014. Automatic denoising of functional MRI data: combining independent component analysis and hierarchical fusion of classifiers. *Neuroimage* 90, 449–468.
- Sato, T.K., Nauhaus, I., Carandini, M., jul 2012. Traveling waves in visual cortex. *Neuron* 75 (2), 218–229.
- Schreiber, T., Schmitz, A., 2000. Surrogate time series. *Phys. Nonlinear Phenom.* 142, 346–382.
- Siegel, J.S., Snyder, A.Z., Ramsey, L., Shulman, G.L., Corbetta, M., 2016. The effects of hemodynamic lag on functional connectivity and behavior after stroke. *J. Cereb. Blood Flow Metab.* 36 (12), 2162–2176.
- Smith, S.M., Beckmann, C.F., Andersson, J., Auerbach, E.J., Bijsterbosch, J., Douaud, G., Duff, E., Feinberg, D.A., Griffanti, L., Harms, M.P., Kelly, M., Laumann, T., Miller, K.L., Moeller, S., Petersen, S., Power, J., Salimi-Khorshidi, G., Snyder, A.Z., Vu, An T., Woolrich, M.W., Xu, J., Yacoub, E., Ugurbil, K., Van Essen, D.C., Glasser, M.F., 2013. Resting-state fMRI in the human connectome project. *Neuroimage* 80, 144–168.
- Smith, S.M., Miller, K.L., Moeller, S., Xu, J., Auerbach, E.J., Woolrich, M.W., Beckmann, C.F., Jenkinson, M., Andersson, J., Glasser, M.F., Van Essen, D.C., Feinberg, D.A., Yacoub, E.S., Ugurbil, K., 2012. Temporally-independent functional modes of spontaneous brain activity. *Proc. Natl. Acad. Sci.* 109 (8), 3131–3136.
- Takahashi, K., Kim, S., Coleman, T.P., Brown, K.A., Suminski, A.J., Best, M.D., Hatsopoulos, N.G., jan 2015. Large-scale spatiotemporal spike patterning consistent with wave propagation in motor cortex. *Nat. Commun.* 6 (May), 7169.
- Thompson, G.J., Pan, W.J., Magnuson, M.E., Jaeger, D., Keilholz, S.D., 2014. Quasi-periodic patterns (QPP): large-scale dynamics in resting state fMRI that correlate with local infraslow electrical activity. *Neuroimage* 84, 1018–1031.
- Townsend, R.G., Solomon, S.S., Chen, S.C., Pietersen, a.N.J., Martin, P.R., Solomon, S.G., Gong, P., mar 2015. Emergence of complex wave patterns in primate cerebral cortex. *J. Neurosci.* 35 (11), 4657–4662.
- Vanni, M.P., Chan, A.W., Balbi, M., Silasi, G., Murphy, T.H., 2017. Mesoscale mapping of mouse cortex reveals frequency-dependent cycling between distinct macroscale functional modules. *J. Neurosci.* 37 (31), 7513–7533.
- Vidaurre, D., Smith, S.M., Woolrich, M.W., November 28, 2017. Brain network dynamics are hierarchically organized in time. *Proc. Natl. Acad. Sci.* 114 (48), 12827–12832.

- Wu, J., Dehkharghani, S., Nahab, F., Allen, J., Qiu, D., 2018. The effects of acetazolamide on the evaluation of cerebral hemodynamics and functional connectivity using blood oxygen level-dependent MR imaging in patients with chronic steno-occlusive disease of the anterior circulation. *AJNR. Am. J. Neuroradiol.* 14 (5), 487–502. <https://doi.org/10.1177/1073858408317066>.
- Wu, J.-Y., Huang, X., Zhang, C., oct 2008. Propagating waves of activity in the neocortex: what they are, what they do. *Neurosci* 14 (5), 487–502.
- Yeo, B.T.T., Krienen, F.M., Chee, M.W.L., Buckner, R.L., 2014. Estimates of segregation and overlap of functional connectivity networks in the human cerebral cortex. *Neuroimage* 88, 212–227.
- Zanos, T.P., Mineault, P.J., Nasiotis, K.T., Guitton, D., Pack, C.C., jan 2015. A sensorimotor role for traveling waves in primate visual cortex. *Neuron* 85 (3), 615–627.
- Zhang, H., Jacobs, J., sep 2015. Traveling theta waves in the human Hippocampus. *J. Neurosci.* 35 (36), 12477–12487.