

## Modelling subject variability in the spatial and temporal characteristics of functional modes

Samuel J. Harrison<sup>a,b,c,\*</sup>, Janine D. Bijsterbosch<sup>a,d</sup>, Andrew R. Segerdahl<sup>a</sup>, Sean P. Fitzgibbon<sup>a</sup>, Seyedeh-Rezvan Farahibozorg<sup>a</sup>, Eugene P. Duff<sup>a,e</sup>, Stephen M. Smith<sup>a</sup>, Mark W. Woolrich<sup>a,b</sup>

<sup>a</sup> FMRI, Wellcome Centre for Integrative Neuroimaging, University of Oxford, Oxford, UK

<sup>b</sup> OHBA, Wellcome Centre for Integrative Neuroimaging, University of Oxford, Oxford, UK

<sup>c</sup> Translational Neuromodeling Unit, University of Zurich & ETH Zurich, Zurich, Switzerland

<sup>d</sup> Department of Radiology, Washington University Medical School, Saint Louis, USA

<sup>e</sup> Department of Paediatrics, University of Oxford, Oxford, UK

### A B S T R A C T

Recent work has highlighted the scale and ubiquity of subject variability in observations from functional MRI data (fMRI). Furthermore, it is highly likely that errors in the estimation of either the spatial presentation of, or the coupling between, functional regions can confound cross-subject analyses, making accurate and unbiased representations of functional data essential for interpreting any downstream analyses. Here, we extend the framework of probabilistic functional modes (PFMs) (Harrison et al., 2015) to capture cross-subject variability not only in the mode spatial maps, but also in the functional coupling between modes and in mode amplitudes. A new implementation of the inference now also allows for the analysis of modern, large-scale data sets, and the combined inference and analysis package, PROFUMO, is available from [git.fmrib.ox.ac.uk/samh/profumo](http://git.fmrib.ox.ac.uk/samh/profumo). A new implementation of the inference now also allows for the analysis of modern, large-scale data sets. Using simulated data, resting-state data from 1000 subjects collected as part of the Human Connectome Project (Van Essen et al., 2013), and an analysis of 14 subjects in a variety of continuous task-states (Kieliba et al., 2019), we demonstrate how PFMs are able to capture, within a single model, a rich description of how the spatio-temporal structure of resting-state fMRI activity varies across subjects.

We also compare the new PFM model to the well established independent component analysis with dual regression (ICA-DR) pipeline. This reveals that, under PFM assumptions, much more of the (behaviorally relevant) cross-subject variability in fMRI activity should be attributed to the variability in spatial maps, and that, after accounting for this, functional coupling between modes primarily reflects current cognitive state. This has fundamental implications for the interpretation of cross-sectional studies of functional connectivity that do not capture cross-subject variability to the same extent as PFMs.

### 1. Introduction

One of the key changes to the landscape of the analysis of functional connectivity via rfMRI in recent years has been the proliferation of large population-level studies (Bamberg et al., 2015; Breteler et al., 2014; Miller et al., 2016; Van Essen et al., 2012b) and multi-site data-sharing initiatives (Biswal et al., 2010; Gorgolewski et al., 2017; Kennedy et al., 2016; Mennes et al., 2013; Poldrack et al., 2013; Scott et al., 2011; Thompson et al., 2014)<sup>1</sup>. This has allowed investigations into the population-level correlates of fine-grained changes in functional connectivity (Allen et al., 2011; Dubois and Adolphs, 2016), with several studies already finding strong links with a variety of behavioural, genetic and lifestyle factors (Colclough et al., 2017; Elliott et al., 2018; Finn et al., 2015; Smith et al., 2015); together, these findings augur well for the search for clinically relevant, personalised predictions from functional neuroimaging data (Abraham et al., 2017; Dubois and Adolphs, 2016; Insel and Cuthbert, 2015; Stephan et al., 2017). In sum, there has

been a shift in what is required of analysis techniques, namely that they must be both interpretable and sensitive to subject-level variability, and at the same time they need to scale to meet the computational demands posed by large data sets.

#### 1.1. Implications of variability over subjects

In this paper, we are primarily interested in the interpretation of—and characterisation of the subject variability in—static functional connectivity<sup>2</sup>. Ultimately, static functional connectivity is encapsulated by the dense connectome—by which we mean the time-averaged voxels-by-voxels connectivity matrix, as defined by the statistical relationships between time courses as extracted from functional data (Friston, 2011;

<sup>2</sup> For the rest of this paper, when we use the term variability in relation to functional measures, it can be assumed to relate to variability in static functional connectivity over subjects or sessions. This does not consider, for example, the moment-to-moment fluctuations characterised as dynamic functional connectivity (Calhoun et al., 2014; Hutchison et al., 2013; Preti and Van De Ville, 2017). This encodes important within-subject state changes (Tagliazucchi and Laufs, 2014), and there is growing evidence that this captures between-subject trait differences too (Vidaurre et al., 2017).

\* Corresponding author.

E-mail address: [samuel.harrison@ndcn.ox.ac.uk](mailto:samuel.harrison@ndcn.ox.ac.uk) (S.J. Harrison).

<sup>1</sup> For a more complete overview of data sharing initiatives, see the NeuroImage special issues (Eickhoff et al., 2016).

Friston et al., 1993). However, dense connectomes are cumbersome computationally, and the natural spatial scale of the functional data is likely to be much lower than the several hundred thousand voxels present in a typical fMRI acquisition (Van Essen et al., 2012a). In practice, what we are seeking is a parsimonious summary of the static functional connectivity that is both readily interpretable and captures key forms of variability.

The canonical approach for analyses of static functional connectivity is to summarise the high-dimensional data in terms of a comparatively small number of either parcels or functional systems<sup>3</sup>. These are usually defined in terms of their spatial configuration, at which point it is possible to extract representative time courses from functional data and analyse these. There will naturally be variability in functional connectivity in several domains, though based on the above framework we will focus on two key ones here: firstly, we will refer to variability in the size, shape and location of functional regions as *subject variability in spatial organisation*; secondly, we will use *subject variability in temporal features* to denote the changes in summary measures based on said time courses—in particular, the strength of functional connectivity between regions (i.e. functional connectomes). Finally, note that for clarity we will use the term *functional coupling* to specifically refer to the functional connectivity between regions as described by these low-dimensional connectomes<sup>4</sup>.

The assumption that is implicit in either the parcel or system-level analyses is that registration to a common space means that the time courses we extract based on group-level spatial descriptions are an accurate, or at least unbiased, description of each subject's data. However, given that it is by no means uncommon to observe three-fold variation in the areal extent of regions of primary visual cortex across subjects (Andrews et al., 1997; Dougherty et al., 2003); or that non-homeomorphic morphological changes, such as subjects exhibiting different number of gyri and sulci, are prevalent (Amiez and Petrides, 2014; Shackman et al., 2011) even in identical twins (Bartley et al., 1997; Hasan et al., 2011); or that macroscale anatomical features are poor predictors of cytoarchitectonic borders (Amunts et al., 2007); then we should expect there to be substantial disparities in the presentation of functionally homologous regions across subjects, even after nonlinear registration (Brett et al., 2002; Devlin and Poldrack, 2007; Mueller et al., 2013; Van Essen and Dierker, 2007). Recent observations have confirmed this for functional data, where it has been shown that this subject variability in spatial organisation 'can give rise to divergent connectivity estimates from the same seed region in different subjects' (Gordon et al., 2017a)—with the results from several studies also suggesting that reorganisations of functionally homologous regions that cannot be represented by diffeomorphic warps seem to be commonplace (Braga and Buckner, 2017; Glasser et al., 2016a; Gordon et al., 2016; 2017b; Hacker et al., 2013; Harrison et al., 2015; Kong et al., 2018; Laumann et al., 2015). Furthermore, these differences have a substantial impact on the data: cross-subject differences in static functional connectivity have been shown to be much larger than either cross-site effects (Noble et al., 2017) or cross-condition, within-subject changes (Gratton et al., 2018).

Loosely speaking, these spatial differences in functional connectivity after registration can arise for four reasons: there will naturally be some errors in the registration process, resulting in structural features that are not brought into correspondence; there will be locations where anatomical landmarks bear little relation to functional subdivisions, meaning structural similarity is not a sufficient condition for accurate registration; there will be genuine non-homotopic reorganisations, whereby the standard registration approaches based on diffeomorphic warps could

never succeed<sup>5</sup>; and there will be dynamic—either moment-to-moment or state-dependent—changes in the functional connectivity structure (Buckner et al., 2013; Krienen et al., 2014; Salehi et al., 2020). If these different sources of variability in spatial organisation are not accounted for, then one expects the inferred mode time courses to be a farrago of contributions from the underlying 'true' set of modes (Allen et al., 2012; Smith et al., 2011). Worse still, if the structural differences capture meaningful cross-subject differences—which they almost certainly will do (Llera et al., 2019)—then the amount of misalignment, and hence the quality of the extracted time courses, will reflect information that is anatomical rather than functional in origin (Bijsterbosch et al., 2018). This breaks the central tenet of investigations into subject variability in temporal features, as we can no longer assume that a group-level description of the functional architecture is a reliable description of individual subjects, or even that we can use these to extract unbiased estimates of functional coupling. How then, do we proceed from here?

The first approach we could take is to improve the registrations, and hope that better algorithms and utilising a richer feature set to drive the alignment will push individual subjects ever closer towards the group description (Robinson et al., 2018; 2014; Tong et al., 2017). Notably however, the multiple recent observations that single functional regions can be manifested as multiple disjoint regions in some subjects, is something that not even advanced functional registration algorithms reliant on diffeomorphic warps can correct for. The minimum requirement for this approach is therefore the use of advanced registration techniques that can non-homotopically reorganise the spatial layout of functional regions, as, for example, introduced by Conroy et al. (2013); Guntupalli et al. (2018, 2016), or Langs et al. (2010).

The alternative approach, and the one that we take in this paper, is to build algorithms that can extract estimates of subject variability in temporal features while simultaneously accounting for the variable presentation of functional regions at the subject level. Several methods have been proposed to do exactly this, using both hierarchical models of functional systems (Abraham et al., 2013; Harrison et al., 2015; Li et al., 2017; Varoquaux et al., 2011) and parcels (Kong et al., 2018; Langs et al., 2016; Liu et al., 2012). We provide a more fulsome description of these, and their counterparts that extract subject-specific information given a fixed group template, in Appendix A.1. However, the majority of these methods have what is potentially a major limitation: the flow of information is almost exclusively from group to subject. In other words, there are only relatively rudimentary efforts to tap into what we might hope is a virtuous cycle: we should be able to use our group-level estimates to infer accurate subject-level information, but, crucially, we should also be able to utilise the observed variability at the subject level to refine our group-level parameterisations. Furthermore, the same process should hold within subjects, such that accurate estimation of the individual spatial presentations should improve evaluation of the temporal information, and vice versa.

Finally, while we have tended to focus on connectomes as the principal temporal feature of interest in the above discussion, there are other types of variability we are interested in. Recent work has shown that, for example, amplitudes—by which we mean any metric which represents the amount of fluctuation in activity of a functional region over time—carry a substantial amount of information about subjects (Bijsterbosch et al., 2017; Duff et al., 2008; Miller et al., 2016; Zang et al., 2007; Zou et al., 2008), provided we are sufficiently careful in how we distinguish changes from those in functional coupling (Duff et al., 2018), and then how we interpret said changes (Qing and Gong, 2016).

<sup>5</sup> It is somewhat contentious whether (structural) registration should be held responsible for the latter two processes. Our definition of registration is somewhat broader, as we hold it responsible for bringing subjects into structural and functional correspondence. While structural registration is unlikely to be sufficient here, this is nevertheless a reasonable aim for multi-modal registration approaches. For a good discussion of these issues see e.g. Van Essen and Dierker (2007).

<sup>3</sup> Resting-state networks, intrinsic connectivity networks, etc.

<sup>4</sup> We make this distinction as the spatial maps, which characterise the location of functional regions, also capture aspects functional connectivity and organisation.

Amplitudes are therefore another type of subject-specific information that we would hope analysis methods could identify, and more importantly disambiguate from, the types of subject variability we have already discussed. This is an illustrative example of the complexity of the task of characterising functional connectivity: at every level of any perceptual hierarchy of features we impose (i.e. separation into spatial and temporal features, or subdivision of temporal features into amplitudes and coupling), we expect there to be multiple ways to identify the different features, and substantial cross-subject variability that is correlated across the different categories.

## 1.2. Outline

For the rest of this paper, we will outline our approach for simultaneously inferring group- and subject-level descriptions of functional systems. We use the term *mode* to describe our mathematical description of a given system.

To begin with, we present our probabilistic model for these modes, including the way we parameterise subject variability in both spatial and temporal features, and our approach for inference. This is a significant extension of the proof-of-concept method (Harrison et al., 2015) in several key ways: we introduce a new hierarchical model to better capture the functional coupling between modes, incorporate a model for mode amplitudes to engender a cleaner separation between different types of functional variability, and we overhaul the entire implementation to help the inference scale to large data sets.

We then compare the performance of our method with existing approaches. We do this using both simulated and empirical data, namely the complete set of rfMRI data as released by the Human Connectome Project and “active-state” fMRI data from a more conventionally sized study. Finally, we offer some brief discussions as to the significance of our results.

## 2. Model

Our approach infers subject-level probabilistic functional modes (PFMs)—each of which can be thought of as being described by a subject-specific spatial map and a set of time courses—across the whole cohort simultaneously. Ensuring that there is correspondence between the inferred modes across the cohort is a challenge (Esposito et al., 2005), especially on resting-state data where we cannot assume any common temporal structure.

However, we can use the information at the group-level to inform the subject-specific decompositions: both the subject-specific spatial maps and the low-dimensional, between-mode functional connectomes are constrained to vary around their group-level descriptions, and we can also leverage the expected properties of the haemodynamic response to further constrain the time courses. Moreover, we can use the subject-specific modes to learn about the variability of all these properties, thereby allowing us to not only describe typical patterns of activity, but to also quantify the extent to which observed patterns are atypical. We do this by building, and then inferring upon, a hierarchical probabilistic model for rfMRI data as described by a set of modes, and it is this that we outline in the following section.

### 2.1. Matrix factorisation models

Defining a mode in terms of a spatial map and time course means that it is fundamentally a matrix factorisation approach, a mathematical formulation which underpins principal component analysis, independent component analysis, non-negative matrix factorisation, dictionary learning and several other of the well established methods for extracting modes from rfMRI data. For completeness, we briefly introduce our notation for this class of models before introducing our extensions.

Firstly, each subject,  $s$ , from a cohort of  $S$  subjects, is scanned  $R_s$  times. Note that we do not assume that each of the runs for a given

subject (i.e.  $r \in \{1, \dots, R_s\}$ ) are identical from a modelling standpoint: they could, for example, represent different time points in a longitudinal study, or different conditions<sup>6</sup>, and we may therefore want to treat them differently. The fMRI data are acquired in  $V$  voxels and at  $T$  time points, which we reshape into a data matrix  $\mathbf{D}^{(sr)} \in \mathbb{R}^{V \times T}$ . We do all our analyses after the data has been registered into a common space, so the number of voxels is constant across subjects. We do however allow the number of time points per run to vary (i.e.  $\mathbf{D}^{(sr)} \in \mathbb{R}^{V \times T^{(sr)}}$ ), but for notational simplicity we drop any superscripts on  $T$ .

The problem we are faced with is defining an extension to the standard matrix factorisation approach to account for these multiple data. In the spatial domain, as discussed in the Introduction, we expect between-subject variability in the locations of functional regions, even after registration, and we expect these effects to be amongst the dominant sources of functional variability. We make the pragmatic decision to focus on differences in the static configuration of functional systems specifically, and we target our spatial approach towards what are essentially misalignments.

Therefore, as in Harrison et al. (2015), we model subject and run variability within the matrix factorisation framework as follows. We are looking for a set of  $M$  modes, and we assume that the subject variability in spatial organisation we observe across subjects, by virtue of it being driven primarily by cortical reorganisations, is consistent across all runs for a given subject. This gives a set of subject-specific spatial maps,  $\mathbf{P}^{(s)} \in \mathbb{R}^{V \times M}$ , that will potentially be observed multiple times. Furthermore, each run will have its own unique set of time courses,  $\mathbf{A}^{(sr)} \in \mathbb{R}^{M \times T}$ , as well as a set of mode amplitudes,  $\mathbf{h}^{(sr)} \in \mathbb{R}^M$ . For convenience we adopt the following convention:  $\mathbf{H}^{(sr)} \in \mathbb{R}^{M \times M} \equiv \text{diag}(\mathbf{h}^{(sr)})$ . Finally, note that in general we infer a small number of PFMs relative to  $V$  and  $T$ , which gives a parsimonious description of the data. However, this means that the factorisation will not be exact, so we express the data as the contribution from the PFMs and a noise term,  $\boldsymbol{\epsilon}^{(sr)} \in \mathbb{R}^{V \times T}$ . This set of assumptions allows us to describe the complete model for one run as

$$\mathbf{D}^{(sr)} = \mathbf{P}^{(s)} \mathbf{H}^{(sr)} \mathbf{A}^{(sr)} + \boldsymbol{\epsilon}^{(sr)} \quad (1)$$

In the following sections, we describe how we model the dependencies between these run-specific decompositions, as well as the key properties of rfMRI data that we are trying to capture. For reference, a full graphical model is provided in the Supplementary Material.

### 2.2. Spatial model

The spatial model remains conceptually similar to the approach we used in Harrison et al. (2015). For each mode, there is a rich group-level description capturing the mean group maps and typical subject variability around these; as Van Essen and Dierker (2007) discuss, in light of subject variability, it is essential that ‘[regions are] represented probabilistically whenever possible, in a way that reflects variability in cortical convolutions and in [their] size, location, and internal (e.g., topographic) organization’. Similarly, subject maps are parameterised such that they retain the key characteristics of the group maps, but allow for genuine variability while being robust to spurious correlations induced by noise.

A key modification we make to the previous model is to change the way we model the spatial map distribution, by relaxing the delta-Gaussian mixture model to a double-Gaussian mixture model. Previously, the weights in voxels which were inferred to be outside of a given mode were set to exactly zero. In reality however, essentially all voxels will exhibit a *weak* correlation with a given mode time course<sup>7</sup>, and, particularly in studies like the Human Connectome Project with thousands

<sup>6</sup> Eyes-open, eyes-closed, pre/post an intervention, various “active-state” paradigms etc.

<sup>7</sup> Cf. the noisy estimates of beta values from a GLM fit.

of time points per subject, there is often sufficient evidence *a posteriori* to model this noise as small, but nevertheless non-zero, weights<sup>8</sup>. The new model allows for exactly this type of ‘spurious’ (i.e. statistically but not biologically significant) correlation by including a noise distribution to capture small deviations from zero in the spatial map weights. While we are not interested in these small weights per se, if we do not include a more explicit noise model then the model will erroneously include them as signal thereby hindering our ability to detect genuine ‘neural’ signal.

This contamination by noise happens for three main reasons. Firstly, as [Bright and Murphy \(2015\)](#) recently showed, even well-characterised functional modes can be identified from noise processes like subject motion. Conversely, this implies that even accurately identified modes may well correlate with non-neural processes. Secondly, given the complex, long-range spatial autocorrelations present in fMRI data ([Kriegeskorte et al., 2008](#)), fMRI noise processes have a non-trivial structure. This is heightened by spatial smoothing, which is an often used pre-processing step for fMRI data (though less so for modern high spatial and temporal resolution data ([Glasser et al., 2016b](#))). This is advantageous as it ameliorates the problem of residual spatial mis-alignment after registration, but induces heightened spatial correlations in the noise. While it would be possible to model this, estimating—and then correcting for—the true number of spatial degrees of freedom in the data is notoriously difficult ([Eklund et al., 2016](#); [Worsley et al., 1996](#)), and would be computationally expensive over a large number of voxels. Finally, in the section on the noise model itself, we demonstrate how unstructured noise can have a stabilising effect on matrix factorisation models. Therefore, we make the pragmatic decision to account for these effects in the spatial model, rather than trying to incorporate a more complex mechanistic model for the noise.

The resulting model takes the following form. For voxel  $v$  in mode  $m$ , the subject-specific spatial weights are distributed as follows:

$$\begin{aligned} p(\mathbf{P}_{vm}^{(s)} | q_{vm}^{(s)} = 1) &= \mathcal{N}(\mathbf{P}_{vm}^{(s)} | \mu_{vm}, \sigma_{vm}^2) \\ p(\mathbf{P}_{vm}^{(s)} | q_{vm}^{(s)} = 0) &= \mathcal{N}(\mathbf{P}_{vm}^{(s)} | 0, (\eta_m^{(s)})^2 \zeta_v^2) \\ p(q_{vm}^{(s)}) &= (\pi_{vm})^{q_{vm}^{(s)}} (1 - \pi_{vm})^{1 - q_{vm}^{(s)}} \\ p(\eta_m^{(s)}) &= \mathcal{N}(\eta_m^{(s)} | 0, \gamma_\eta^2) \end{aligned} \quad (2)$$

Where  $q_{vm}^{(s)}$  is a binary indicator variable which represents whether a given voxel’s weight is drawn from the signal or the noise component.

This distribution is defined in terms of several group-level hyperparameters: the probability that a given weight is drawn from the signal rather than the noise distribution,  $\pi_{vm}$ ; the mean and standard deviation of the signal component,  $\mu_{vm}$  and  $\sigma_{vm}$  respectively; and the new parameters, the standard deviation of the noise component, which we parameterise as  $\eta_m^{(s)} \zeta_v$  for reasons which we explain in detail later.

Note how much richer this description is than the single set of group-level means that most currently used techniques infer. For example, the  $\sigma_{vm}$  parameters can capture the types of spatial non-uniformity in subject variability observed by [Mueller et al. \(2013\)](#). Therefore, when inferring subject maps, the inference will automatically be informed by the data more than the group mean in regions inferred to exhibit high functional heterogeneity over subjects, and vice versa for regions with low subject-to-subject variability.

The model also includes the set of distributions over the group-level hyperpriors (see the Supplementary Material for the way these, and all subsequent, hyperparameters are specified). Starting with the hyperpriors on the ‘signal’ component, we place a mixture model prior over the group means, which, as in the previous work, is inspired by the spike-slab distribution ([George and McCulloch, 1993](#); [Ishwaran and Rao, 2005](#); [Mitchell and Beauchamp, 1988](#); [Titsias and Lázaro-Gredilla, 2011](#)). This encourages sparsity in the group-level spatial maps, thereby

<sup>8</sup> See [Colclough et al. \(2018\)](#) for a discussion of exactly this effect in relation to inference of functional couplings between regions.

encoding ideas about functional segregation, as well as allowing more flexibility when specifying the distribution of the non-zero weights. However, we introduce an extension and model the non-zero weights with a combination of two Gaussians with different variances. This allows the group-level distribution of non-zero spatial weights to have heavier tails than the single Gaussian used in the previous incarnation of the model.

$$\begin{aligned} p(\mu_{vm} | \rho_{vm} = 2) &= \mathcal{N}(\mu_{vm} | \tau_{\mu 2}, \gamma_{\mu 2}^2) \\ p(\mu_{vm} | \rho_{vm} = 1) &= \mathcal{N}(\mu_{vm} | \tau_{\mu 1}, \gamma_{\mu 1}^2) \\ p(\mu_{vm} | \rho_{vm} = 0) &= \delta(\mu_{vm}) \\ p(\rho_{vm}) &= \prod_{i \in \{0,1,2\}} (\lambda_{\mu i})^{[\rho_{vm} = i]} \end{aligned} \quad (3)$$

Where  $\rho_{vm}$  is the probability that a voxel in the group map is drawn from each of the three distributions, and  $[\rho_{vm} = i]$  is the Iverson bracket.

The group signal standard deviations,  $\sigma_{vm}$ , take an inverse-gamma hyperprior:

$$p(\sigma_{vm}) = \Gamma(\sigma_{vm}^{-2} | a_\sigma, b_\sigma) \quad (4)$$

Returning to the hyperpriors on the ‘noise’ component, in [Eq. 2](#), the standard deviation of the noise component of the subject-specific spatial map distribution is parameterised as  $\eta_m^{(s)} \zeta_v$ . The  $\zeta_v$  parameter encodes spatial inhomogeneity in the noise variance: for example, we expect more structured noise due to motion around the edges of the brain; similarly, we expect more physiological noise in the brainstem. This group noise standard deviation,  $\zeta_v$ , also takes an inverse-gamma hyperprior:

$$p(\zeta_v) = \Gamma(\zeta_v^{-2} | a_\zeta, b_\zeta) \quad (5)$$

However, we also expect different signal-to-noise ratios, both across subjects and modes. Therefore, we include an extra parameter,  $\eta_m^{(s)}$ , which captures variations in the noise level<sup>9</sup>. We place a weak prior on  $\eta_m^{(s)}$ , as we want the overall scale of each spatial map to be determined by the signal rather than the noise, as this makes cross-subject analyses more informative:

$$p(\eta_m^{(s)}) = \mathcal{N}(\eta_m^{(s)} | 0, \gamma_\eta^2) \quad (6)$$

Finally, the last hyperprior to specify is that on the group membership probabilities. This follows a beta distribution:

$$p(\pi_{vm}) = \beta(\pi_{vm} | a_\pi, b_\pi) \quad (7)$$

In summary, the model has rich descriptions of the spatial maps, both at the group and subject level, and allows us to encode typical patterns of variability. Furthermore, while we have included a weak sparsity constraint at the group-level, there is no explicit constraint on, for example, orthogonality of the spatial maps. Therefore, the model can capture modes that are highly spatially overlapping in what is arguably a more natural way than independent component analysis—even despite a historic tendency to overstate those criticisms ([Beckmann et al., 2005](#); [Calhoun et al., 2013](#); [Smith et al., 2012](#)).

One last point to note is that when we present our results, the group maps we show are the marginal posterior means of the whole spatial distribution, rather than the  $\mu$  parameters themselves. The group-level maps are therefore  $E[\pi_{vm} \mu_{vm} | D]$ , which has the nice property that it incorporates the uncertainty about whether each voxel is drawn from the signal or the noise component.

### 2.3. Temporal model

In the temporal domain, the unconstrained nature of rfMRI data means that we can say relatively little about the time courses from a given run, as there are no external events from which we can search

<sup>9</sup> See e.g. [Gelman \(2006\)](#) for a related discussion of redundant parameterisations of variance.



for consistent time-locked patterns of mode activation. However, functional connectomics has shown that, as well as having a consistent group structure, both the interactions between modes and simple amplitude measures encode interesting information about subjects. Similarly, the haemodynamic processes lend neural processes a distinct temporal signature. That being the case, we wish to formulate a model that primarily captures these two phenomena.

However, we expect the inferred time courses to be corrupted by noise, even if we properly make allowances for the global noise process  $\epsilon^{(sr)}$ . As mentioned in the Spatial Model section, there are likely to be structured noise processes that violate our haemodynamic assumptions. This needs to be accounted for before we can introduce the targeted models of the BOLD signal.

Analogously to the spatial model, we extend the model from Harrison et al. (2015) by making the pragmatic decision to allow noisy time courses. Therefore, our time course model contains two terms: the first represents the clean BOLD time courses,  $\mathbf{B}^{(sr)}$ , while the second represents the noise that corrupts these,  $\xi^{(sr)}$ . This gives:

$$\mathbf{A}^{(sr)} = \mathbf{B}^{(sr)} + \xi^{(sr)} \quad (8)$$

There is an additional benefit of this explicit parameterisation of the BOLD time courses. Recent work has claimed that the [fractional] Amplitude of Low Frequency Fluctuations ([f]ALFF) (Zang et al., 2007; Zou et al., 2008; Zuo et al., 2010a), as derived from rfMRI data, captures aspects of subject variability related to disease. Our parameterisation allows us to derive a related quantity, which we term the fractional amplitude of BOLD time courses (fABT). This is simply defined as the power in the clean BOLD time courses  $\mathbf{B}^{(sr)}$ , relative to the power in the noise time courses  $\xi^{(sr)}$ , calculated for each mode and each run individually. Conceptually, this is very closely related to fALFF, but it has the clear advantage that it does not require defining 'low' frequencies in terms of an arbitrary threshold; rather, the signal of interest is based on an explicit model of the HRF. Secondly, the calculated fABT measures specifically relate to the activity in different functional systems which makes the measure more interpretable.

### 2.3.1. Haemodynamic model

We use the haemodynamic response function (HRF) based model that we introduced in Harrison et al. (2015). This is a relatively simple, computationally efficient, linear model that captures the gross properties of the HRF via the temporal autocorrelations that it induces in the data. We assume a white noise 'neuronal' process convolved with a canonical HRF<sup>10</sup>, whose autocorrelation function we can capture using a full covariance matrix,  $\mathbf{K}_B \in \mathbb{R}^{T \times T}$ , for all the time points in a given run. As the overall variance of the time courses is arbitrary given the explicit amplitude parameters, we simply ensure that  $\mathbf{K}_B$  is scaled such that all entries on the main diagonal are unity.

### 2.3.2. Subject-level mode interactions

The major extension relative to the previous model is an explicit parameterisation of the functional coupling between modes. As discussed earlier, we expect to observe temporal interactions between modes, and this will lend some structure to the mode time courses. We define these interactions in terms of the precision matrix between the mode time courses. In other words, we combine the HRF-derived autocorrelation structure with a prior on the between-mode precision matrix,  $\alpha^{(sr)} \in \mathbb{R}^{M \times M}$ , in a matrix normal distribution.

The combined prior on the haemodynamic time course for all the PFM's in a given run is then:

$$p(\mathbf{B}^{(sr)} | \alpha^{(sr)}) = \mathcal{MN}(\mathbf{B}^{(sr)} | \mathbf{0}, \alpha^{(sr)^{-1}}, \mathbf{K}_B) \quad (9)$$

<sup>10</sup> For adult populations, both the SPM double-gamma HRF (Friston et al., 2007) or the principal component of the FLOBS basis in FSL (Woolrich et al., 2004) are provided, though this can be replaced for different populations as appropriate e.g. (Arichi et al., 2012).

### 2.3.3. Group-level mode interactions

The temporal interactions between modes have been characterised as having a consistent structure across the group (Shehzad et al., 2009), so we introduce a hierarchical model to capture this. Subject- or run-level variability will manifest itself as deviations from this set of group interactions. This formulation we use is, in essence, the same model as that proposed by Marrelec et al. (2006), but where we have two principal advantages: firstly, inference is informed by the full posteriors on the rest of the model (i.e. rather than point estimates); and, secondly, that the regularisation that arises from these priors will inform the inference of the rest of the model parameters.

Starting at the subject level, we estimate the subject/run-specific temporal precision matrix  $\alpha^{(sr)}$  to keep track of the functional connectivity between modes. These precision matrices follow a Wishart distribution, and we introduce a hyperparameter,  $\beta \in \mathbb{R}^{M \times M}$ , that encourages the interactions to be consistent across subjects and/or runs. This takes the form of a hyperprior on the subject-specific scale matrices, and again this follows a Wishart distribution.

$$p(\alpha^{(sr)} | \beta) = \mathcal{W}(\alpha^{(sr)} | a_\alpha, \beta) \quad p(\beta) = \mathcal{W}(\beta | a_\beta, \mathbf{B}_\beta) \quad (10)$$

Furthermore, we can also place restrictions on the type of variability we want the model to capture. If, for example, subjects are scanned multiple times but always under the same conditions, then it may well be appropriate to generate a consensus set of interactions for that subject by pooling over runs. We can do this straightforwardly by setting  $\alpha^{(sr)} \equiv \alpha^{(s)}$ . Alternatively, if the runs vary across the group in a consistent way (e.g. 'before' and 'after' scans) then we may want to explicitly model these conditions as separate entities. We can do this by introducing a family of group-level interactions,  $\{\beta^{(r)}\}_{r=1}^R$ , and selectively using these as the hyperpriors on  $\alpha^{(sr)}$  as appropriate. This gives us enormous flexibility and allows us to increase our statistical power by making targeted assumptions about the key modes of variation.

### 2.3.4. Time course specific noise model

The noise time course of mode  $m$  at time  $t$ ,  $\xi_{mt}^{(sr)}$ , is simply drawn from a Gaussian distribution with precision  $\omega_m^{(sr)}$ . This gives

$$p(\xi_{mt}^{(sr)} | \omega_m^{(sr)}) = \mathcal{N}(\xi_{mt}^{(sr)} | 0, \omega_m^{(sr)^{-1}}) \quad (11)$$

Each  $\omega_m^{(sr)}$  takes a gamma hyperprior:

$$p(\omega_m^{(sr)}) = \Gamma(\omega_m^{(sr)} | a_\omega, b_\omega) \quad (12)$$

### 2.4. Amplitude model

Again, the amplitude model is an extension to our previous work. This has a straightforward formulation, with these parameters simply designed to account for the run-to-run variations in the overall activity of each mode. These are parameterised in terms of  $\mathbf{H}^{(sr)} \equiv \text{diag}(\mathbf{h}^{(sr)})$ , and follow a Gaussian distribution:

$$p(\mathbf{h}_m^{(sr)} | \mu_h, \Sigma_h) = \mathcal{N}(\mathbf{h}_m^{(sr)} | \mu_h, \Sigma_h) \quad (13)$$

The group-level parameters,  $\mu_h$  and  $\Sigma_h$  capture any consistent cross-subject relationships between the mode amplitudes. For example, Bijsterbosch et al. (2017) recently reported that the amplitudes of sensorimotor modes are correlated with one another, as are the amplitudes of cognitive networks. It is exactly these types of effects that these hyperpriors are able to capture.

The hyperpriors are formulated as follows:

$$p(\mu_h) = \prod_{m=1}^M \mathcal{N}((\mu_h)_m | \tau_{\mu_h}, \gamma_{\mu_h}^2) \quad (14)$$

$$p(\Sigma_h) = \mathcal{W}(\Sigma_h^{-1} | a_h, \mathbf{B}_h) \quad (15)$$

Furthermore, we impose a post-hoc positivity constraint on these variables as part of the inference procedure. As there is a multiplicative ambiguity as to the signs of the components in a matrix factorisation model, we can do this without loss of generality.

## 2.5. Noise model

The final part of the model left to specify is the noise process,  $\epsilon^{(sr)}$ , which we assume is zero-mean, white Gaussian noise, with an overall precision for each run,  $\psi^{(sr)}$ . This specifies the likelihood:

$$\begin{aligned} p(\epsilon^{(sr)}) &= \mathcal{MN}(\epsilon^{(sr)} | \mathbf{0}, (\psi^{(sr)})^{-1} \mathbf{I}_V, \mathbf{I}_T) \\ &= p(\mathbf{D}^{(sr)} - \mathbf{P}^{(s)} \mathbf{H}^{(sr)} \mathbf{A}^{(sr)}) \end{aligned} \quad (16)$$

This noise precision then takes a standard gamma hyperprior:

$$p(\psi^{(sr)}) = \Gamma(\psi^{(sr)} | a_\psi, b_\psi) \quad (17)$$

This relatively simple structure assumes that the noise variance is the same in every voxel, which is particularly useful as it allows us to exploit the properties of the matrix normal distribution, leading to very computationally efficient inference (Stegle et al., 2011). We can preprocess the data in such a way that this is a reasonable assumption to make, and this is discussed in Appendix A.2.

What is perhaps more problematic is that this model does not acknowledge the spatial smoothness of fMRI data, which means that the noise is not truly independent over voxels. It would be possible to model this, for example by inferring a full spatial covariance matrix for the noise that acknowledged the dependencies between voxels that smoothing introduces. Again, we decide that the benefits of this more complex model are outweighed by the increased computational burden, and again we discuss a way in which we can mitigate the effects of this model misspecification via a relatively straightforward adjustment for the spatial degrees of freedom introduced by Groves et al. (2011), as discussed in Appendix A.4.

### 2.5.1. Spatially and temporally specific noise models

One of the key changes to the model as introduced here and its previous incarnation is the way we model noise on the spatial maps and timecourses, as well as the overall noise described above. Interestingly, these different sources of noise can be beneficial for matrix factorisation models even in the absence of the fMRI-specific effects we postulated.

To demonstrate this, we use a simple, single-run version of our generative model,  $\mathbf{D} = \mathbf{P}\mathbf{A} + \epsilon$ , and we assume the maps and timecourses are full rank to simplify the derivations below. The ordinary-least-squares single-regression estimator for the spatial maps,  $\hat{\mathbf{P}}_{[sr]}$ , given the ground-truth timecourses is:

$$\hat{\mathbf{P}}_{[sr]} = \mathbf{D}\mathbf{A}^{-1} = \mathbf{P} + \epsilon\mathbf{A}^{-1} \quad (18)$$

If we instead run dual regression—using the Woodbury matrix identity for the key rearrangements—we find a different estimator for  $\hat{\mathbf{P}}_{[dr]}$ :

$$\begin{aligned} \hat{\mathbf{A}}_{[dr]} &= \mathbf{P}^{-1}\mathbf{D} = \mathbf{A} + \mathbf{P}^{-1}\epsilon \\ \hat{\mathbf{P}}_{[dr]} &= \mathbf{D}\hat{\mathbf{A}}_{[dr]}^{-1} \\ &= (\mathbf{P} + \epsilon\mathbf{A}^{-1}) - (\mathbf{P}\mathbf{A} + \epsilon)\mathbf{A}^{-1}\mathbf{P}^{-1}(\mathbf{I} + \epsilon\mathbf{A}^{-1}\mathbf{P}^{-1})^{-1}\epsilon\mathbf{A}^{-1} \\ &= (\mathbf{P} + \epsilon\mathbf{A}^{-1}) - \epsilon\mathbf{A}^{-1} \\ &= \mathbf{P} \end{aligned} \quad (19)$$

What is surprising is that the dual regression estimator is closer to the ground truth, even though the intermediate timecourses,  $\hat{\mathbf{A}}_{[dr]}$ , are noisy. This unintuitive behaviour occurs because dual regression involves two regressions on the same noise, and this has concrete implications for the PFM model. When we fit the haemodynamic model to the timecourses, we exclude the temporally specific noise terms from the estimation of the functional coupling between modes. However, we need to include the temporal noise terms when using the timecourses to estimate the subject-specific spatial maps, as removing it could increase the variance of the inferred maps. The situation is directly analogous with the model for spatial noise: while it is not a quantity of interest for

cross-subject modelling, its inclusion can improve the stability of the overall estimation.

In sum, the PROFUMO approach uses the spatially and temporally specific noise where the stabilising effect on matrix factorisation models means that it is expedient to do so, but seeks to avoid letting it confound cross-subject analyses. By way of contrast, dual regression is not naturally able to separate these types of noise.

## 2.6. Inference approach

We use a computationally efficient variational Bayesian approach to infer upon the probabilistic model outlined above. This technique is well established for graphical models that have a conjugate-exponential structure, as is the mean-field approximation that renders the inference tractable (Attias, 2000; Blei et al., 2017; MacKay, 2003; Winn and Bishop, 2005); as such, we will not cover the details of that here. In the Appendices, we outline several of the implementation details, including our data preprocessing pipeline, the way we handle large data sets, tweaks to the model and the initialisation procedure.

The combined inference and analysis package, PROFUMO (from PRObabilistic FUNctional MOdes) is available from [git.fmrib.ox.ac.uk/samh/profumo](https://git.fmrib.ox.ac.uk/samh/profumo) and is compatible with FSL (Jenkinson et al., 2012). All subsequent analyses were performed with version 0.11.1.

The model clearly has a large number of hyperparameters, but as described in the Supplementary Material we can drastically reduce the effective number given that the overall variance of the data is fixed by the internal preprocessing. Furthermore, the vast majority of the parameters that need setting govern the group-level hyperpriors and, as such, are several steps removed from the subject-level decompositions. This means that we can use the same default hyperpriors for all the analyses presented here, and that the inference generalises well across simulated, volumetric, and surface-based data, as well as datasets with very different numbers of subjects.

## 2.7. Model summary


In summary, we explicitly model many of the properties of rfMRI data within the PROFUMO framework. In the spatial domain, we have a complex group-level model that captures both mean effects and typical patterns of variability, and use these to regularise the subject-specific spatial maps. The temporal model is based around the physiological properties of the BOLD signal, and includes another hierarchical model for the functional coupling between modes. Similarly, we capture differences in the overall activity levels of modes via the amplitude parameters. Finally, we can generate additional summaries by combining parameters as desired, which includes, for example, the measures related to the fractional amplitudes of the BOLD signal.

## 3. Results

Here, we demonstrate the performance of PROFUMO using a set of simulated data and two empirical datasets. All comparisons are with spatial independent component analysis and dual regression (ICA-DR) (Beckmann et al., 2005; Calhoun et al., 2001; Nickerson et al., 2017; Zuo et al., 2010b), as this is what has been used in previous publications on the empirical data.

### 3.1. Simulations

The simulation framework is explicitly designed to be challenging, such that it tests the various ways in which the assumptions the different models make are most likely to be violated. This includes spatial and temporal correlations between components; spatial variability, including a model for misalignments; amplitude variability across subjects and components; a (weakly) nonlinear HRF that varies over both



**Fig. 1.** Performance of different algorithms on simulated data. For each metric, optimal performance is shown by the horizontal green line. The metrics are accuracy in recovery of the subject-specific spatial maps, recovery of the run-specific network matrices (netmats), recovery of cross-subject differences in amplitudes (as different approaches normalise the data differently, we look at relative changes in amplitudes across subjects), and any biases in the recovered temporal correlations towards the spatial correlation structure. As well as PROFUMO and ICA-DR, we test dual regression starting with the group-level ground-truth spatial maps (GTg-DR) and thresholded dual regression (ICA-DRt, GTg-DRt).

subjects and space; and spatial and temporal smoothness in the residuals. This extends previously published analyses (Bijsterbosch et al., 2019; Harrison et al., 2015), and all simulation code is available from [git.fmrib.ox.ac.uk/samh/PFM\\_Simulations](https://git.fmrib.ox.ac.uk/samh/PFM_Simulations).

Specifically, we simulate data containing 15 components in a group of 40 subjects, each with two runs containing 10,000 voxels and 500 timepoints at a TR of 2.0s. A more detailed overview of the data generation procedure is provided in Appendix A.6. We then test how well PROFUMO and ICA-DR can recover the ground-truth parameters, pooling results across 10 different simulated datasets. Finally, to give more detailed insights into the performance of ICA-DR we include several intermediate steps: firstly, to separate the performance of ICA and dual regression, we include a dual regression analysis starting from the group-level ground-truth spatial maps (GTg-DR); secondly, we include the thresholded variant of dual regression proposed by Bijsterbosch et al. (2019) which is designed to reduce the observed bias in functional coupling (ICA-DRt, GTg-DRt).

Four key performance metrics are shown in Fig. 1, and a much more detailed set of comparisons is included in the Supplementary Material. PROFUMO is able to accurately recover spatial maps, amplitudes and functional coupling network matrices (netmats), and much more so than either ICA-DR or the improved thresholded variant (ICA-DRt).

Crucially, the inferred PFMs are also unbiased in the presence of spatio-temporal correlations between components, unlike ICA-DR. What Bijsterbosch et al. (2019) demonstrated was that inaccurate estimation of the group-level spatial correlation structure—an inevitable consequence of the orthogonality constraints of ICA—leads to biased estimates of functional coupling. What we show here is a stronger result: this effect is present even when starting from the correct group-level spatial maps (GTg-DR). In this case, the effect is driven by the mismatch between the true subject-level spatial correlations and those between


the group-level maps. In other words, this bias will be present for all dual regression analyses, however the group-level maps are generated.

Furthermore, in the Supplementary Material, we repeat the simulations but with the addition of structured noise, including subject-specific artefacts that can be either spatially specific or global. While the differences between methods are less pronounced, there are still clear benefits to using PROFUMO. However, performance does suffer, and, as such, we strongly recommend ICA-based artefact removal before running PROFUMO, as is the case for the two empirical datasets presented here.

### 3.2. Human connectome project data

To evaluate the ability of PROFUMO to detect subtle subject-specific variations in functional connectivity, we use data from the Human Connectome Project (HCP) (Van Essen et al., 2013; 2012b). This is for two main reasons. Firstly, the most recent data release includes high-quality functional data from over 1000 subjects and, as such, is an ideal test for methods that purport to be suitable for population-level studies as mentioned in the Introduction. Secondly, the functional pipeline has been published (Smith et al., 2013a) and the results are available to download—thereby offering a comparison that is independently verifiable. The pipeline uses spatial ICA and dual regression to characterise subject variability in both spatial and temporal features. While it would also be possible to examine the equivalent pipeline based on temporal ICA, this has not been used so extensively: for example, the HCP’s MegaTrawl analyses are based on the spatial ICA pipeline<sup>11</sup>. Similarly, this pipeline does not make use of the new thresholded variant of dual regression.

<sup>11</sup> <https://db.humanconnectome.org/megatrawl/index.html>.



**Fig. 2.** Group-level spatial maps for four example PFMs, as inferred from the HCP data. The PFMs are (a) the default mode network (DMN) (Buckner et al., 2008; Greicius et al., 2003; Raichle et al., 2001; Shulman et al., 1997); (b) a mode described as a variant of the DMN by Braga and Buckner (2017); (c) a mode with strong spatial anticorrelations with the DMN; and (d) the mode containing functional activity within POS2 (Glasser and Van Essen, 2011).

Based on the simulated data, this would improve the results slightly, though PROFUMO still outperforms this variant on essentially all of the metrics we tested. Again, the aim is to use the existing, publicly available results as a baseline.

A key aim of modern, large-scale studies of functional connectivity is to relate neurobiological changes to individual differences in genetic, lifestyle and behavioural factors. Using the HCP data also allows us to do this by comparing our results with a wide range of information about subjects. The data involves a battery of cognitive tests, and also records a range of metrics based on health and lifestyle: we will refer to differences in these as *subject variability in behavioural measures*. We can indirectly assess the effects of genetics and environment by calculating the heritability of key imaging metrics; we do this by utilising the fact that many twins and siblings were involved in the study. Finally, we can examine *subject variability in structural measures* by relating functional measures to the thicknesses, areas and volumes of key cortical and subcortical structures as derived from the structural MRI scans (Glasser et al., 2013). In this way, we can quantify to what extent different methods are able to capture key aspects of functional variability, and if there are meaningful relationships with other measures.

A more detailed overview of the data, and the tests we carry out here, can be found in Appendix A.7.

### 3.2.1. Analyses

Both PROFUMO and spatial ICA were run at a dimensionality of 50, at which point the modes were reordered for visualisation and noise components—or, in the case of PROFUMO, modes eliminated by the Bayesian model complexity penalties—were removed. Even on the extensive and high-quality HCP data, PROFUMO does not identify more than 50 PFMs: when run at higher dimensionalities, more PFMs are simply eliminated from the model. We discuss why PROFUMO is likely to be conservative in this regard in more detail later.

For the full HCP data, PROFUMO therefore infers the posterior over approximately 25,000,000,000 parameters (1000 subjects, 100,000 grayordinates, 50 modes, 5 parameters per grayordinate). In terms of computational requirements, this analysis took approximately 110 hours using 18 cores on a single compute node, and memory usage peaked at 350GB.

Finally, note that subsequent figures display spatial maps on the cortical surface for simplicity and concision. However, all grayordinates

(comprising approximately 60,000 cortical vertices and 30,000 subcortical voxels (Glasser et al., 2013)) were used in all analyses.

### 3.2.2. Overview of the PFM spatial model


To begin with, in Figs. 2 and 3 we show examples of the group- and subject-level spatial maps for four PFMs in order to demonstrate the richness of information contained within the PFM model. We do this to emphasise that PROFUMO is able to identify PFMs with strong spatial relationships with one another (in terms of overlap and anticorrelations), while at the same time being able to identify complex, subject-specific reorganisations of the group templates.

The most striking feature of the subject maps in Fig. 3 is simply how much variability relative to the group maps there is. These results are from data already aligned using surface-based registration driven by functional features, which arguably represent the current 'gold-standard'; for warp-based registrations (Coalson et al., 2018; Glasser et al., 2016b). Despite this, and as we and several others have demonstrated, there are pronounced differences between subjects, with both shifts in the relative location of functional regions over surprisingly large distances, and complex, non-homotopic splittings and reorganisations of the regions themselves. Furthermore, as highlighted in the figure, even though the PFM itself is large, there are several subject-specific features that are too small to be accurately represented at the typical spatial scale of parcellations applied to fMRI.


However, while the descriptions of modes in terms of the mean group- or subject-level spatial maps are familiar, a key advantage of the PFM framework is the more detailed group-level parameterisation. In other words, we can go beyond simply noting the degree of subject variability: we can now quantify it in detail on a per-mode level. In Fig. 4 we again take the default mode network [Figs. 2 and 3] as an example and plot the four key group-level spatial parameters: the probability that a given voxel belongs to the DMN, the mean and variability over subjects of the signal component of the DMN's voxelwise weights, and the standard deviation of the spatial noise component. The information encoded by the mean weights is familiar, but the other parameters add novel and complementary information.

For example, the memberships [Panel (b)] demonstrate that default mode activity is distributed over a surprisingly large area, with consistently detected activity across much of the lateral prefrontal cortex. This is an effect that has been captured by several recent, high-powered





**Fig. 3.** Subject-level equivalents of the default mode network shown in Panel (a) of Fig. 2.




**Fig. 4.** Example of the key group-level spatial parameters for the PFM representing the default mode network [Panel (a) of Fig. 2 and Fig. 3], as inferred from the HCP data. The parameters are the (a) posterior means of the signal component,  $\mu_{vm}$ ; (b) posterior memberships,  $\pi_{vm}$ ; (c) posterior standard deviations of the signal component,  $\sigma_{vm}$ ; (d) posterior standard deviations of the noise component,  $\zeta_v$ .

single-subject analyses (Gonzalez-Castillo et al., 2012; Huth et al., 2016; Laumann et al., 2015; Poldrack et al., 2015). However, while the activity is widespread, it is also distinct: the areas of high and low probability are sharply delineated. Similarly, the standard deviations [Panel (c)] add extra information by telling us about variability in the size of the weights—that is, in the strength of the detected activity—and we can see that, in this instance, the activity in the inferior parietal lobule is much more variable in strength across subjects than that in the precuneus.

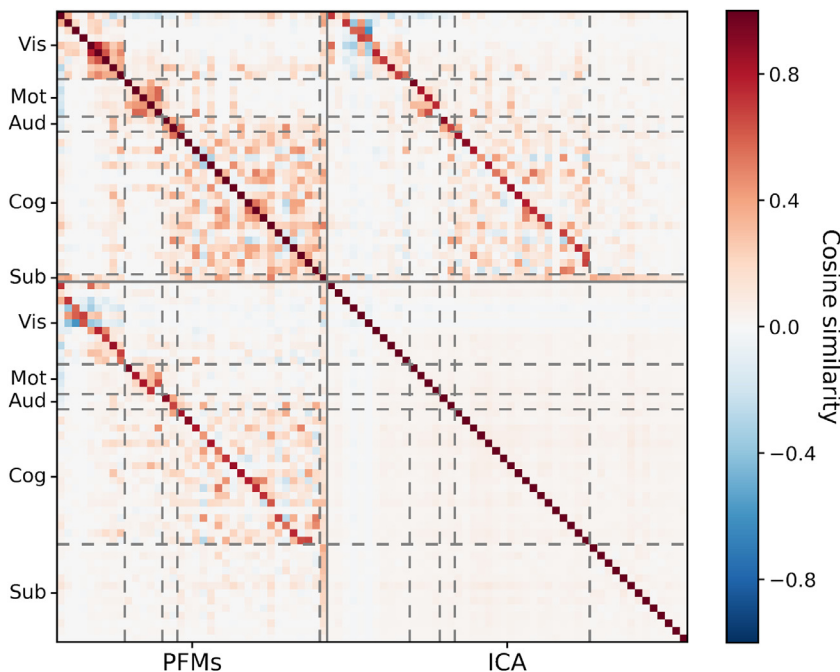
This detailed characterisation of non-homogeneous variability across the cortex is a key advantage of the more complex group-level model we have adopted, and we expand upon this in Fig. 5. This summarises the membership probabilities and weight standard deviations across all modes. There is a clear pattern whereby association cortex contains more overlapping modes than sensory cortices [Panel (a)], and that the spatial weights are also more variable in association cor-

tex [Panel (b)]—note how this is in agreement with the results of Mueller et al. (2013). Finally, the uncertainty in the memberships themselves [Panel (c)] tells us about shifts in locations between subjects. For example, note the very clear area of variability in medial frontal cortex between SMA and pre-SMA (Johansen-Berg et al., 2004). This metric is presumably particularly sensitive to this region because variability here tends to manifest itself as relatively simple anterior-posterior shifts of the SMA/pre-SMA boundary, whereas more complicated 2D rearrangements of overlapping PFMs are present elsewhere.

In summary, the PFM spatial model captures familiar group-level modes, and exhibits many of the complex subject-specific rearrangements already described in the literature. However, the key advantage is the way in which we have parameterised this model. Crucially, the richness of the group description allows us to make specific claims about the patterns of variability across the population that are ordinarily hard to tease apart.



**Fig. 5.** Summaries of the group-level spatial parameters encoding different aspects of variability across subjects. The panels are (a) mode overlap; (b) variability in mode strength; (c) variability in mode memberships. Mode overlap is defined as the posterior memberships averaged across all modes,  $\frac{1}{M} \sum_{vm} \pi_{vm}$ . Variability in mode strength is captured by the weighted average of the posterior standard deviations,  $(\sum_{vm} \pi_{vm} \sigma_{vm}) / (\sum_{vm} \pi_{vm})$ . Finally, variability in mode memberships is given by the average entropy, in bits, of the membership distributions,  $\frac{1}{M} \sum_{vm} H(\pi_{vm})$ .



**Fig. 6.** Spatial similarity between the sets of group-level spatial maps as inferred by PROFUMO and ICA. Modes were split into five categories and reordered: visual (Vis); motor (Mot); auditory (Aud); cognitive (Cog); and subcortical (Sub). This ordering is used for all subsequent sections.

### 3.2.3. Comparison with spatial ICA

To begin with, we examine the performance of the different models in terms of their inference of the group-level spatial descriptions. In Fig. 6 we plot the similarity between these group-level descriptions.


There are several key points to note. Firstly, there are strong spatial correlations between the PFM maps, especially within the different categories. By way of contrast, the independence assumptions in spatial ICA preclude this. Secondly, PROFUMO is relatively conservative: it only infers 36 signal modes compared to the 48 found by ICA, and the difference is particularly pronounced in the subcortical regions. This subcortical difference is predominantly driven by the different signal properties of the HCP data between cortical and subcortical grayordinates, and the different data normalisation strategies the two algorithms use. The result is that ICA tends to find subcortical regions appearing in components without much cortical involvement, whereas PROFUMO tends to find subcortical regions appearing in components with cortical involvement. Finally, despite the above differences, there is fundamentally a strong relationship between the two sets of maps. Most cortical modes appear in both decompositions, and often look fairly similar; this is encouraging, as we do not expect a radically different patterns of functional connectivity at the group level given how many published methods have converged on similar descriptions.

### 3.2.4. Properties of subject variability in spatial organisation


Given that the group-level descriptions are fairly similar between PFMs and ICA, the obvious question are to what extent does the extra group-level information in the PFM model regularise the subject-specific decompositions, and in what ways do the subject-specific maps diverge from the group-level representations? We deal with the former first, and in Fig. 7 we look at that the consistency of the subject maps as inferred by PROFUMO and the ICA-DR pipeline. As expected given the regularisation from the more complex group-level priors, the PFM maps are much more consistent across subjects.

However, this increase in consistency could also be explained if the subject-specific PFM spatial maps were simply pushed closer to the group maps by the priors, thereby being less faithful to the 'true' patterns of functional connectivity at the subject level. While this does not appear to be the case for the exemplar subject maps, what we really want to quantify is whether they are capturing 'interesting' aspects of subject variability in spatial organisation. In other words, are the differences between the approaches meaningful, and do they make different predictions about the subjects themselves?


To investigate this, we use the fact that the HCP includes data from twins and siblings to investigate the influence of genetics and environment. We estimate the voxelwise broad-sense heritability of the subject-



**Fig. 7.** Similarity between the subject-specific spatial maps, for both PFMs and ICA-DR, as inferred from the HCP data. For each voxel and in every pair of subjects, we compute the Pearson correlation coefficient between the two  $M$ -dimensional vectors of mode weights. The maps shown here are the correlation coefficients averaged over every pair of subjects.



**Fig. 8.** Analyses of the heritability of the subject-specific mode maps, for both PFMs and ICA-DR, as inferred from the HCP data. In (a) and (b) we display the voxelwise estimates of broad-sense heritability ( $H_b^2$ ), and in (c) we compare the two as a scatter plot.



specific spatial maps we observe: in each voxel and each subject, we extract the vector of PFM or ICA map weights, and look to see if these weight vectors are more consistent in monozygotic than dizygotic twins (see Appendix A.7 for full methodological details). The results of this analysis are shown in Fig. 8.


The results show a clear increase in heritability for the PFM spatial maps, suggesting that they are more sensitive to subject variability that we can attribute to genetic factors. Furthermore, this is not simply attributable to a reduction in noise or as the result of the priors pushing the subject maps closer to the group. While the PFM maps are more consistent across subjects than ICA-DR [Fig. 7], the heritability relates to the difference in consistency between monozygotic and dizygotic twins and, as such, a global increase in consistency is not enough to explain the increased heritability.

We can also gain further insights into this observation by utilising the HCP's retest data. 46 subjects underwent the full HCP imaging and behavioural testing protocol twice, of which there is full fMRI data from 42. This allows us to examine how the algorithms perform on the hitherto unseen retest scans. The group-level representations from the


full data (i.e. the ICA spatial maps, and the group-level PFM posteriors) were used to derive new subject maps from the independently acquired retest data.

In Fig. 9 we compare subject-specific realisations of the language mode as derived by PROFUMO and the ICA-DR pipeline. This particular mode was chosen because a characteristic split in area 55b in some subjects was reported and examined in some detail by Glasser et al. (2016a). In terms of a comparison between PROFUMO and ICA-DR, both are clearly sensitive to the same gross re-organisations that occur. For example, both can detect the rearrangement of area 55b in the original and retest data for the subject shown here. However, the most marked difference is in the noise-level and appearance of anticorrelations. Relative to ICA-DR, the PFMs show much reduced background noise in regions not associated with the networks, and do not exhibit anticorrelations (indicated by negative weights, shown in blue) tightly interposed between positive weights. This is presumably a simple consequence of dual regression's inability to separate signal from noise, as we discussed in the section on noise modelling. By way of contrast, the information encoded by the group-level parameters in the PFM model suppresses the back-





**Fig. 9.** Example spatial maps for the language mode, for both PFMs and ICA-DR, as inferred from the full HCP data and the HCP retest data for subject 149337. Only the left lateral surface is shown.



**Fig. 10.** Specificity of the subject-level spatial maps as inferred from both the original and retest HCP data by PROFUMO and ICA-DR. The group results from the full data are used to derive subject-specific spatial maps in the unseen retest data. In (a) we show the similarity of the inferred maps in the same subject, separately for each mode. In (b) we calculate the fingerprint specificity, or how much more similar the maps in the same subject are as compared to maps from non-matching pairs of subjects, averaged over modes. This is equivalent to the difference between the diagonal and the off-diagonal elements (calculated for each column separately) in the full similarity matrices as shown in (c) and (d).

ground noise in regions that are not part of the language network, but in a way that does not preclude inferring complicated rearrangements of functional regions.

To assess the reliability of the different decompositions on the retest data more quantitatively, we assess the specificity of the inferred spatial maps as ‘fingerprints’ that uniquely identify different subjects (Finn et al., 2015; Horien et al., 2019). This is shown in Fig. 10.

Firstly, we compute the spatial similarity between the new subject-specific spatial maps from the retest data, and the original set from the full data, for every pair of subjects. We pool these retest results over all modes and subjects, and this is shown in Panel (a). Again, the subject-specific PFM maps are much more consistent across the two acquisitions.


Secondly, we assess whether this leads to more specific fingerprints. In Panel (b) we show that the fingerprint specificity (i.e. the amount by which the two sets of maps from the same subject are more similar than paired maps from different subjects) is also higher for the PFMs. In other words, not only are the maps generally more consistent across subjects, but there is an increase in subject specificity above and beyond this effect.

In summary, the comparisons with ICA-DR have demonstrated that while the group-level descriptions are similar, the more complex hierarchical modelling in PROFUMO allows us to infer spatial maps that are more consistent—on both the original data and the held-out retest data—as well as being more specific and capturing more informative aspects of cross-subject variability.

### 3.2.5. Overview of the PFM temporal model

Here, we briefly give a summary of the key temporal parameters—that is, the amplitudes and the functional coupling between modes—as inferred by PROFUMO on the HCP data. Note again that these are new parameters: in other words, it was only possible to investigate these in a post-hoc fashion based on the previous PFM model. Firstly, in Fig. 11 we plot the cross-subject correlations between the mode amplitudes, as captured by the  $\Sigma_h$  parameter. Encouragingly, we see a clear replication of the results of Bijsterbosch et al. (2017), who reported strong correlations between the amplitudes of sensorimotor modes, as well as between cognitive modes, but relatively weak correlations across the two categories. However, the crucial difference be-





**Fig. 11.** Cross-subject relationships between the amplitudes of the PFMs, as inferred from the HCP data. For visualisation purposes, we display the posterior precision matrix,  $\Sigma_h$ , after transforming it to both full and partial correlation coefficients.

tween this result and the original observation is that this behaviour was initially demonstrated from a purely post-hoc analysis of the ICA-DR results, whereas it is explicitly parameterised and inferred within the PFMs model. What this means is that this knowledge of the systematic relationship between mode amplitudes is available *during* inference, and it is therefore naturally incorporated as an extra factor regularising the subject-specific decompositions.

Secondly, in Fig. 12 we plot the PFM functional coupling parameters,  $\beta$  and  $\alpha^{(s)}$  (these represent the group- and subject-level temporal network matrices respectively). What is striking is how weak the functional coupling is between modes in the group-level network matrix (netmat), especially given that we have an explicit hierarchical model to allow for just these interactions. This is not trivial to explain away as a spatial effect either: despite the fact that these interactions are more similar to what we would expect from temporal ICA, the PFM spatial maps are similar to those inferred by spatial ICA which typically infers strong functional coupling between modes. We quantify the implications of this different view on functional coupling from the PFM model in the following section.

### 3.2.6. Multivariate relationships with behavioural variables

How then, are we to interpret the differences between the PFM and ICA-DR approaches? Do they simply represent a different trade-off between sensitivity and specificity in the spatial and temporal domains, or are they telling us something fundamentally different about brain activity?


To probe this further, we performed a series of multivariate analyses to investigate the different ways in which the two models encode cross-subject information. Like in Smith et al. (2015), canonical correlation analysis (CCA)—a multivariate analysis technique used to find the linear relationships between sets of variables (Hotelling, 1936)—was used to summarise the key correspondences (see Appendix A.7 for methodological details). Furthermore, as some sets exhibit more than one strong linear relationship, we use the RV coefficient (Robert and Escoufier, 1976)

to give a principled summary of the multivariate information reported by the CCA. In Fig. 13, we examine the full set of pairwise relationships between the behavioural and structural variables from the HCP, and the spatial maps, amplitudes and network matrices from both PROFUMO and ICA-DR.

There are several key results we can glean from this analysis. Firstly, the cross-subject information captured by the different aspects of the PFM model is relatively distinct. Comparing the similarity between the PFM measures with those for the ICA-DR variables (i.e. the on-diagonal blocks), we can see that the scores are typically lower for the PFMs. In other words, the temporal measures derived from the PFMs carry relatively different information from the spatial measures about the subjects themselves, at least compared to their ICA-DR equivalents.

Secondly, if we examine the relationships with the behavioural and structural measures in the bar graph on the right, there are several striking differences between the methods. As we would expect from our earlier analyses, the PFM spatial maps are the best predictors of structural variables. They are also good predictors of the behavioural variables, though slightly less so than the ICA-DR netmats. However, the stories for the temporal information are very different. The PFM amplitudes, fABT and netmats are relatively poor predictors of both behavioural and structural variables, though, intriguingly, they are better predictors of behaviour than structure. By way of contrast, the ICA-DR amplitudes and netmats are better behavioural predictors, though surprisingly they are also good predictors of structure (e.g. one can predict the sizes and thicknesses of cortical areas better than behavioural measures from the ICA-DR amplitudes).

Given the simulation results, the interpretation is relatively straightforward: the ICA-DR pipeline contains inherent biases that conflate spatial and temporal information. Furthermore, even though we do not explicitly test it here, it is interesting to note that using the thresholded version of dual regression to correct this bias also reduces the correlation between temporal netmats and behaviour (Bijsterbosch et al., 2019). In other words, and consistent with the results on simulated data, thresh-



**Fig. 13.** Relationships between the cross-subject information encoded by different analyses. The non-functional variables (NFVs) have been separated into variables from the HCP's battery of behavioural tests, and variables derived from structural MRI relating to brain size and morphology. On the left we plot the log RV coefficient calculated between the subspaces of the top ten CCA components as calculated between every pair of sets of variables, and on the right we reproduce the relationships with the non-functional variables (i.e. the top two rows / two leftmost columns) as a bar chart for ease of visualisation. Higher values of the RV coefficient indicate that more similar cross-subject information is being captured.

olded dual regression is an improvement on ICA-DR but is less accurate than the full PROFUMO model. The question that remains however, is what information, if any, is the PFM temporal model capturing if not the trait-like behavioural variables examined here?

### 3.2.7. Summary

Given the full set of results presented on the HCP data, the implication is that the PFMs, by virtue of the improved spatial modelling in particular, are better able to capture interesting information about cross-subject variability in spatial organisation. However, this does not address the relative lack of information encoded in the various temporal measures that PFMs capture. We address this point using another data set in the following section.

### 3.3. Active-state data

Given the way that subject variability in spatial and temporal features simultaneously co-varies with a wide range of non-imaging derived subject measures, it is very challenging to conclusively disambiguate them from studies like the HCP. However, if we manipulate the functional connectivity at the subject level, for example by changing the cognitive state (Gratton et al., 2018; Kieliba et al., 2019; Krienen et al., 2014; Salehi et al., 2020; Shirer et al., 2011; Vanderwal et al., 2017), then we can begin to examine temporal differences in more detail. Crucially, by looking at multiple conditions for the same subject we essentially eliminate the influence of structural variability from the functional data.

To do this, we use a dataset collected where subjects were scanned when in different active states—these are induced by performing simple, continuous tasks in the scanner, of which rest (i.e. eyes-open fixation) is just one (Duff et al., 2018; Kieliba et al., 2019; Sala-Llonch et al., 2019). There are five runs for every subject, each collected under different steady-state conditions: a standard resting-state acquisition (Rest); a finger-tapping based motor task (Mot); a passive visual condition (Vis); an independent combination of the visual stimulus and motor task (V-M); and a condition where the specifics of the motor task changed based on the visual stimulus (V+M). A more detailed descriptions of the tasks and data itself can be found in Kieliba et al. (2019). Furthermore, this dataset offers a validation of our method on data acquired using a more conventional sequence and scan duration than the HCP, with fewer subjects, shorter scan durations, and all analyses performed on volumetric rather than surface-based data.

#### 3.3.1. Analyses

As per the modelling assumptions, PROFUMO infers one consensus spatial map per subject, but a separate set of time courses per run. We choose to infer run-specific temporal precision matrices,  $\alpha^{(sr)}$ , with a

consistent group-level hyperprior,  $\beta$ , which is shared across all conditions. Note that we could have chosen to use condition-specific group-level priors,  $\{\beta^{(r)}\}_{r=1}^R$ , but this has the side-effect of invalidating the assumptions behind any subject-level statistics where we compare between conditions. In short, it reduces the cross-subject, within-condition variance which invalidates the typical null hypothesis we use. We leave the problem of performing statistical inference on these types of models for future investigations.

We infer 30 modes for both PROFUMO and ICA-DR, which again seems to be close to the upper limit for PROFUMO on this relatively small dataset. Again, artefactual modes were eliminated and those remaining were reordered for visualisation. In terms of computational requirements, the PROFUMO analysis took approximately 12 hours using 15 cores on a single compute node, and memory usage peaked at 25GB. Compared to the HCP analysis, the demands are higher than expected given the number of subjects for two reasons: firstly, the volumetric analysis contains over twice as many voxels as grayordinates; secondly, we do not do within-subject data reduction for this analysis.

For the ICA-DR pipeline, we use MELODIC (Beckmann et al., 2005; Beckmann and Smith, 2004) to infer a set of group maps, followed by dual regression to generate the run-specific time courses.

#### 3.3.2. Overview of the PFM model

In Fig. 14 we show the group-level properties of the default mode as inferred from this data set. This is directly comparable with Fig. 4 and simply demonstrates that we are able to infer similar summaries of the mode itself, and heterogeneous variability, from fourteen subjects rather than one thousand.

In Fig. 15, we demonstrate some of the properties of the inferred time courses from the PFMs. This data is more challenging than the HCP in that the runs are shorter, and the data has not benefited from resampling onto the cortical surface. Nevertheless, the HRF-based prior constraint results in a temporally smooth timecourse, which we are able to cleanly separate from the high-frequency noise which contaminates them. Furthermore, this is stable when we undo the temporal blurring that the HRF induces, with straightforward estimation of the underlying 'neural' process via whitening with respect to the autocorrelation induced by the HRF.

Finally, in Fig. 16, we display examples of the network matrices to illustrate the typical patterns of, and subject variability in, the functional coupling between PFMs. Interestingly, in this data, PROFUMO infers PFMs with much stronger functional coupling between them than at the run level from the HCP data.

#### 3.3.3. Comparison with ICA-DR

One would hope that the PFM model allows us to more accurately infer the true functional coupling between modes. To begin with, we look at the relationships between the condition-specific network matrices as




Fig. 14. Example of the key group-level spatial parameters for the PFM representing the default mode network, as inferred from the active-state data. The parameters are as per Fig. 4, along with the group map. The panels are the (a) group map; (b) posterior means,  $\mu_{vm}$ ; (c) posterior memberships,  $\pi_{vm}$ ; (d) posterior standard deviations,  $\sigma_{vm}$ ; (e) posterior noise standard deviations,  $\zeta_v$ .



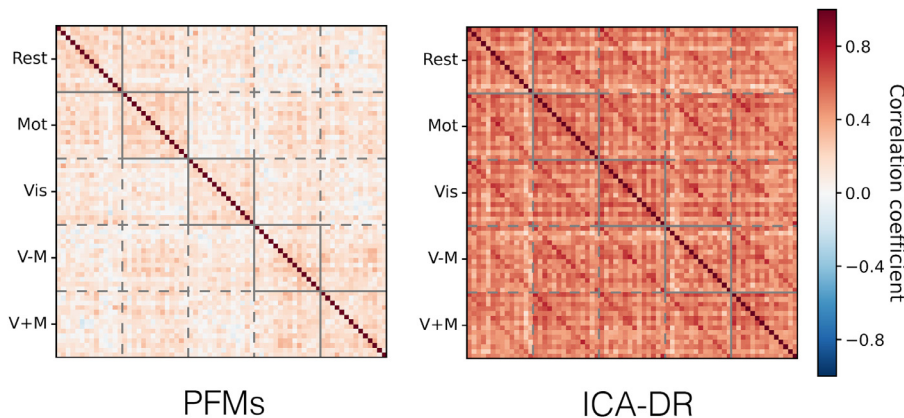


Fig. 15. Example PFM time courses, and observed frequency content, from the active-state data. **Panel (a):** Example time course for one mode in one run. 'Combined' refers to the time course which includes the noise terms ( $A^{(sr)} = B^{(sr)} + \xi^{(sr)}$ ), 'clean' refers to the BOLD portion specifically ( $B^{(sr)}$ ), while 'decorrelated' refers to the clean time course after correcting for the temporal autocorrelation induced by the HRF ( $B^{(sr)} K_B^{-\frac{1}{2}}$ ). **Panels (b) & (c):** Frequency content of the combined and clean time courses respectively, pooled over all runs and subjects. The magnitude of the DFT coefficients are calculated for each time course, and for visualisation purposes, we fit a gamma distribution to the histogram of observed magnitudes for each frequency bin. The mode of this distribution is plotted in red, and the grey region represents the 95% highest density interval (Kruschke, 2014).





**Fig. 16.** Example PFM network matrices, capturing the functional coupling between the mode timeseries. We display the group network matrix alongside the network matrices from subject 25 in the rest, motor and visual conditions. As in Fig. 12, we display the posterior precision matrices (i.e.  $\beta$  for the group level and  $\alpha^{(sr)}$  at the run level) as partial correlations. Modes were split into three categories and reordered for visualisation of the network matrices: visual (Vis); motor (Mot); and cognitive (Cog).



**Fig. 17.** Correlations between the network matrices, for both PFMs and ICA-DR, as inferred from the active-state data. The network matrices are grouped by condition, and the subjects have a consistent ordering within each block. Correlation is the Pearson correlation coefficient between the unwrapped upper-triangle of the network matrices.

inferred by PROFUMO and ICA-DR. These are shown in Fig. 17. While the PFM network matrices are less consistent between conditions and subjects than their ICA-DR counterparts, there is some indication that there is condition-specific modulation across subjects (as indicated by the block diagonal). By way of contrast, the ICA-DR network matrices are dominated by the subjects themselves (i.e. the multiple strong off-diagonal lines in the ICA-DR plot), with no real indication of condition-specific modulations.

In summary, ICA-DR computes netmats that are more similar within subjects than they are within conditions across subjects. By way of contrast, PROFUMO infers netmats that are somewhat more similar within conditions than within subjects. Again, this suggests that the different models for subject variability in spatial organisation have a profound influence on downstream estimates of functional connectivity.

Next, we test whether the different conditions induce focal changes to the between-mode patterns of functional connectivity. The results of a statistical analysis that looks for modulations at the level of individual network matrix edges are shown in Fig. 18. Both the PFM and ICA-DR pipelines detect changes in the coupling of visual regions induced by the visual stimulus, and it appears they both have similar sensitivity to the changes in coupling induced by the changes in cognitive state. There are some differences between the methods: for example, the visual changes detected by PROFUMO are more consistent across the three conditions with visual stimuli than for ICA-DR. Similarly, the types of changes for the combined visuo-motor condition are somewhat different, with ICA-DR finding changes in amplitude predominantly, whereas there are more changes in coupling for PROFUMO.

However, the results are fundamentally fairly similar and the numbers of edges that exhibit significant changes is relatively low—and, perhaps, lower than we might expect given the strong manipulations


of cognitive state<sup>12</sup>—suggesting that the statistical power might be the limiting factor here, especially given that there are only 14 subjects included in this analysis. Finally therefore, we do one further set of tests to probe whether the multivariate information in the network matrices and amplitudes captures condition-specific information. Repeating the analysis of Sala-Llonch et al. (2019), we investigate whether a support vector machine (SVM) can be trained to distinguish between network matrices from different conditions. The accuracy of the SVM classification is tested using a leave-one-subject-out cross-validation framework (Varoquaux et al., 2017), of which we provide more methodological details in Appendix A.8.

As well as comparing PROFUMO and ICA-DR in this way, we additionally examine the effect that the haemodynamic model has on the temporal information that we infer. In other words, can the changes to estimates of functional connectivity be attributed to the advanced spatial modelling alone, or does the regularisation in the time domain improve our estimates too? As well as the explicitly inferred PFM network matrices, we do a post-hoc estimation of the temporal network matrices based on both the BOLD time courses and the combined time courses (i.e.  $A^{(sr)}$ , which includes both the BOLD and noise time courses) to assess what, if any, effect the modelling hierarchy has.

The results from the SVM analysis are presented in Fig. 19. The SVM achieves a significantly better classification accuracy when trained on the PFM netmats, as opposed to those estimated by ICA-DR. Again, this suggests that by correcting for subject variability in spatial organisation the PFM framework allows us to estimate state-induced changes in func-

<sup>12</sup> Note also that Fig. 2 from Sala-Llonch et al. (2019) uses FDR with  $q = 0.2$  for the background, whereas the tests here use a more stringent FWE  $p < .05$  test for significance.





**Fig. 18.** Changes in between-mode functional connectivity as induced by different active states relative to the rest condition. The raw difference between the group mean network matrix during the active condition and during rest is shown above the diagonal, and any significant changes ( $p < 0.05$ ) are highlighted by blue or red squares, for increases and decreases in coupling respectively, below the diagonal. Changes in amplitudes are shown on the diagonal. All tests were family-wise error corrected and computed using the accelerated permutation inference in PALM (Winkler et al., 2016; 2014). The black dots denote elements that were significant under an  $f$ -test over all contrasts. As per Fig. 16, modes were split into three categories and reordered for visualisation of the network matrices: visual (Vis); motor (Mot); and cognitive (Cog).

tional coupling with greater fidelity. By way of contrast, the conflation of spatial and temporal information by ICA-DR masks these more subtle state-related changes in functional coupling. Finally, there appears to be a distinct performance improvement when using the inferred PFM network matrices, suggesting that the hierarchical temporal modelling is advantageous and that we are not discarding relevant information by focusing on the predominantly low-frequency HRF-derived time courses.

#### 4. Discussion

In summary, the results presented above demonstrate three key attributes of PROFUMO. Firstly, the algorithm is applicable to modern, large-scale data, whereby it is exquisitely sensitive to cross-subject variability in spatial organisation. Secondly, the joint inference framework allows estimation of subject variability in temporal features that does not appear to be confounded by spatial differences, which at times leads to a radically different view of functional connectivity. Finally, the implication of these results is that after accounting for spatial variability, the functional coupling between modes is much more reflective of current cognitive state rather than trait-like qualities.

Furthermore, we have shown that there is significant value added in terms of interpretability from the practitioner's point of view in using models of this form. To give a few concrete examples, we can only make the claims pertaining to the dissociation of non-homogeneous spatial variability—as illustrated in Fig. 5—if we can both consistently identify equivalent functional systems across multiple subjects and model the different ways in which variability can arise. Similarly, the ability to capture cross-subject amplitude effects (Fig. 11) or use the model to define alternatives to, for example, fALFF-type measures (Fig. 13) means that many of what would have been post-hoc analyses can be simplified and made more interpretable.

##### 4.1. Group- versus subject-level approaches

The comparisons in this paper have been with ICA-DR, as this is probably the most common method for finding functional modes from resting-state data and is a key part of the HCP's pipelines. How-

ever, while PROFUMO and ICA-DR try and infer on many of the same quantities, they make fundamentally—and not necessarily compatible—assumptions about the data itself.


The key difference between the two is the way PROFUMO entails a holistic model for group- and subject-level representations, whereas ICA-DR assumes they are separately estimable. The majority of group-level ICA methods assume all subjects are in a common space, and proceed to analyse the data without recourse to individual decompositions. This formulation gives much more flexibility for the group-level decomposition to utilise the extra statistical power that concatenating over subjects affords, which means that the ICA modes depart—at times fairly radically—away from what we can resolve at the subject level. As such, ICA seems to be able to identify up to several hundred plausible components, that ultimately begin to resemble a parcellation (Kiviniemi et al., 2009; Smith et al., 2013b).

However, what we show here is that group-level representations are not enough. In the simulated data, even if the ground truth is known at the group-level, the subject-level information inferred by dual regression will be biased and noisy.

What PROFUMO attempts to do is to model as many different facets of multi-subject rfMRI data together as is plausible. Here, we expand on two concrete implications of this approach as compared to other methods.

Firstly, the implication of the joint subject-level modelling in PROFUMO is that for a mode to appear at the group-level it has to be resolvable in the majority of individual subjects. Therefore, this engenders a fundamentally different view on what the dimensionality of the data is. The Bayesian model complexity penalties seem to result in no more than thirty or forty PFMs being identified, essentially regardless of the pre-specified model dimensionality. While more subjects do offer increased regularisation of the subject-level modes, this can only do so much. This is why the inferred number of PFMs is on the same order as the number of signal components as inferred by ICA-FIX ( $23.3 \pm 6.6$  at the run-level for HCP data (Marcus et al., 2013)).

Secondly, we have demonstrated the importance of modelling different characteristics of the data together. In the simulated data, even if the ground truth is known and thresholded dual regression is used



**Fig. 19.** Posterior classification rates for a multi-class SVM trained to distinguish between the different active-state conditions. The results on the left are when the off-diagonal elements of the network matrices are fed in, and the results on the right are when the amplitudes are used as features. Posterior densities are based on the number of correct and incorrect classifications out of the full set of 70 tests (14 subjects; 5 conditions), combined with Haldane's uninformative beta prior (Haldane, 1932). The modes of the distributions are shown by the black bars, and the chance level is shown by the dashed blue line. The two p-values are calculated via McNemar's test (mid-p variant) and Bonferroni corrected (Fagerland et al., 2013). For the PFM netmats, the variants are: *PFMs*: network matrices inferred as part of the PFM model,  $\alpha^{(sr)}$ . *PFM (BT)*: network matrices estimated as the partial correlations between the PFM BOLD time courses  $B^{(sr)}$ . *PFM (CT)*: network matrices estimated as the partial correlations between the combined time courses  $A^{(sr)} = B^{(sr)} + \xi^{(sr)}$ .

to reduce the inherent spatiotemporal biases (Bijsterbosch et al., 2019), PROFUMO is still more accurate than ICA-DR like approaches. Similarly, in the classification of the active-state data, there are clear performance benefits from modelling the netmats hierarchically, even after the spatial variability has been accounted for. This is not to say that the PROFUMO model is perfect, as it clearly contains many simplifying assumptions. However, it is at least an internally consistent framework within which one can begin to explore the implications of different modelling decisions.

#### 4.2. Spatial representations

One of the key messages from this work, in line with other recent reports (Braga and Buckner, 2017; Glasser et al., 2016a; Gordon et al., 2016; 2017a; 2017b; Hacker et al., 2013; Harrison et al., 2015; Kong et al., 2018; Laumann et al., 2015), is that complex rearrangements of functional regions in individual subjects are ubiquitous and of a surprisingly large spatial scale. Figs. 3 and 9 provide reasonable examples of these effects. Even after the advanced multi-modal, surface-based registration employed by the HCP, one often observes spatial rearrangements where subject-specific features are shifted relative to the group by many millimetres.

The difficulty we face when working at the group-level is that the summary features we extract are not necessarily representative of those at the subject-level; they are, and should always be thought of as, probabilistic representations (Van Essen and Dierker, 2007). As discussed in the previous section, we cannot automatically expect that it will be straightforward to project group-level results back to meaningful characterisations of functional connectivity at the subject level. Furthermore, the characteristic size of misalignments probably represents a limit in terms of the size of functional features we can project from the group back to the subject-level; while the native resolution of the subject-level data may well be higher, methods that work on the functional data alone like ICA-DR or PROFUMO will always struggle in the absence of additional constraints if the misalignments are large enough to mean some regions do not overlap with their group-level homologues at all.

In other words, misalignments are now often larger in scale than the fundamental resolution limits imposed by the physics and physiology that governs the properties of the data itself. Subject-level representations are limited by the properties of the data itself: 2mm isotropic voxels are now common, and the spatial characteristics of the HRF do not appear to blur much beyond this (Shmuel et al., 2007); at the group-level, the effective resolution of the data relates to the characteristic size of these residual misalignments between subjects, and these are likely to be larger. What this means is that functional MRI currently occupies an interesting liminal space, where the spatial resolution of high-powered single-subject analyses can now surpass that of studies that employ multitudes of subjects. This probably explains the recent resurgence of exploratory studies based on small numbers of subjects (Braga and Buckner, 2017; Gonzalez-Castillo et al., 2012; Gordon et al., 2017b; Huth et al., 2016; Laumann et al., 2015; Poldrack et al., 2015; Raemaekers et al., 2014; Salehi et al., 2020). Fortunately, recent work has suggested that there is scope to further reduce the size of the residual misalignments (Guntupalli et al., 2018), and use multi-modal data to help identify regions at the subject-level (Glasser et al., 2016a), both of which will be essential parts of the push towards finer spatial scales.

Finally, these observed spatial differences also have implications for parcel-based analyses. Given the many fine-scale variations in the spatial maps and the amount of overlap between PFMs, it may be that we need multivariate analysis techniques that go beyond one summary time course per parcel to capture the richness of the functional data at sub-parcel spatial scales (Anzellotti and Coutanche, 2018; Geerligs et al., 2016; Haak et al., 2018).

#### 4.3. Interpreting spatiotemporal connectivity patterns

One of the striking differences between PROFUMO and ICA-DR is their inferred patterns of functional coupling between regions. Not only do these suggest fundamentally different group-level coupling strengths, but the predictive power at the subject and run level is also different. Whereas ICA-DR netmats primarily correlate with trait-like properties, PROFUMO netmats are more sensitive to changes in cognitive state. Here, we expand on these observations as a final discussion point.

Clearly, there is a complicated relationship linking spatial variability and the functional coupling between modes, and indeed concerns about the interpretability of functional connectivity in the presence of anatomical variability are far from new (Brett et al., 2002). The effect that subject variability in spatial organisation might have on its temporal counterpart has been noted in simulation studies. For example, Allen et al. (2012) observed a sharp decrease in the ability of a variant of ICA-DR to detect subject-specific modes in the presence of subject variability in spatial organisation, an effect which was compounded by spatial overlap between modes<sup>13</sup>. This links to functional coupling via the work of Smith et al. (2011), who noted that if ROIs were misspecified such that the time courses contained a range of contributions from

<sup>13</sup> Figures 4 and 5 in particular.

the true underlying regions, then '[t]he results are extremely bad'. It is the latter result in particular which is particularly shocking: if we do not extract accurate subject-level estimates of functional regions then it is essentially impossible to characterise the functional coupling between them.

Furthermore, a key claim of the related recent work on subject variability in functional connectivity by [Bijsterbosch et al. \(2018\)](#) is that it is not possible to make meaningful claims about what drives cross-subject changes in functional coupling between regions if said regions are not properly delineated at the subject level. In other words, spatial variability does not simply make it harder to estimate functional coupling, it can also fundamentally bias our inferences. Again, the results here—particularly the simulations—extend these results, showing that the way in which dual regression biases functional connectivity estimates away from the spatial correlation structure ([Bijsterbosch et al., 2019](#)) is really an inherent property of mapping between group and subject levels in this way. While this bias can be reduced with the thresholded variant of dual regression, the simulation results, and short theoretical analysis on the role of noise, suggest that the PFM model will be much more performant than this variant.

What we show here with regards to the predictive power of the PROFUMO netmats is that, in line with other work ([Bijsterbosch et al., 2019; 2018; Pervaiz et al., 2020](#)), they are relatively poor predictors of trait-like quantities. Instead, we have shown that they are much more predictive of current cognitive state. However, for analyses that try to use functional coupling to make predictions about individual subjects ([Abraham et al., 2017; Dadi et al., 2019; Pervaiz et al., 2020](#)), the ICA-DR netmats are likely to produce more accurate predictions. In that case, one has to contend with the fact that the induced biases reduce the interpretability of the findings, which may or may not be desirable depending on the specifics of the problem at hand ([Stephan et al., 2015; 2017](#)). Of course, the presence of confounds that are themselves behaviourally relevant—such as head motion ([Couvry-Duchesne et al., 2014; Hodgson et al., 2017; Laumann et al., 2017; Power et al., 2012; Satterthwaite et al., 2012; Van Dijk et al., 2012](#)), physiological noise ([Glasser et al., 2018; Power et al., 2017](#)) or brain volume ([Bartley et al., 1997; McDaniel, 2005; Qing and Gong, 2016](#))—makes this problem of interpretability very challenging in practice for any method.

The results we have presented here suggest that the spatial information encoded by PROFUMO is likely to give much better predictive performance in this context. This is similar to other work which has demonstrated increased performance of spatial features such as, for example, task-based maps ([Bijsterbosch et al., 2018](#)) or parcel topography ([Kong et al., 2018](#)), and, furthermore, that this has a close relationship with structural information ([Llera et al., 2019](#)). The obvious questions are therefore why do spatial rearrangements of functional regions seem to be so predictive in cross-subject analyses, and how do we interpret them? One hypothesis is that this variability in spatial organisation of functional regions is simply reflecting variability in the brain's macroscale structure, for which there are already well established links between environmental, genetic and lifestyle factors ([Douaud et al., 2014; Elliott et al., 2018; Noble et al., 2015; Reiss et al., 1996; Shaw et al., 2006; Stein et al., 2012](#)).

However, it would be an enormous surprise if this reductionist reading of these functional changes as simply reflecting structural variability is the whole story, especially after the registration approaches used. Rather, it is vitally important to understand both what mechanisms give rise to these spatial changes, and, in particular, what unique information does the functional variability carry *over and above* what can be derived from other techniques and modalities.

## 5. Conclusions

All analyses of complex, multivariate functional data require us to make simplifying assumptions, and, as such, the results we see are inevitably coloured by the modelling choices we make. This might involve,

for example, deciding whether to run a parcel- or mode-based analysis, or when choosing which specific method to use. As such, it is essentially impossible to conclusively determine whether one method more accurately characterises the general organisational principles or subject variability from the functional data alone. However, we feel that the above results demonstrate that PROFUMO and the PFMs model are providing a novel and worthwhile perspective on the analysis and interpretation of functional MRI data. We hope that this approach—by virtue of having a model tailored to the properties of fMRI data, the enhanced spatial sensitivity and specificity, and the way spatial variability is automatically accounted for when estimating functional coupling—proves useful.

## CRedit authorship contribution statement

**Samuel J. Harrison:** Conceptualization, Methodology, Software, Validation, Formal analysis, Data curation, Writing - original draft, Writing - review & editing, Visualization. **Janine D. Bijsterbosch:** Conceptualization, Software, Validation, Writing - review & editing. **Andrew R. Segerdahl:** Investigation, Data curation, Writing - review & editing. **Sean P. Fitzgibbon:** Software, Validation, Writing - review & editing. **Seyedeh-Rezvan Farahibozorg:** Methodology, Software, Validation, Writing - review & editing. **Eugene P. Duff:** Conceptualization, Investigation, Data curation, Writing - review & editing. **Stephen M. Smith:** Conceptualization, Methodology, Writing - review & editing, Supervision, Project administration, Funding acquisition. **Mark W. Woolrich:** Conceptualization, Methodology, Writing - review & editing, Supervision, Project administration, Funding acquisition.

## Acknowledgements

We would like to offer our profound thanks to Tamar Makin for making the active-state data set available, and to Roser Sala-Llonch and Sasidar Madugular for the preprocessing. We would also like to thank the anonymous reviewers, whose feedback and suggestions led to many improvements to the manuscript. Data were provided [in part] 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. S.J.H., J.D.B. & S.M.S. were funded by the [Wellcome Trust](#) (grants 098369/Z/12/Z and 091509/Z/10/Z). S.J.H. was also supported by the grant #2017-403 of the Strategic Focal Area “Personalized Health and Related Technologies (PHRT)” of the ETH Domain. S.M.S. also received funding from an MRC Mental Health Pathfinder grant (MC\_PC\_17215). A.R.S. received funding for this work from the following sources: National Institute for Health Research Oxford Biomedical Research Centre, Medical Research Council of Great Britain and Northern Ireland, the Wellcome Trust (London, UK) and the Innovative Medicines Initiative Joint Undertaking (Brussels, Belgium), under grant agreement no 115007 resources of which are composed of financial contribution from the European Union's Seventh Framework Programme (FP7/2007–2013) and EFPIA companies' in kind contribution. A.R.S. is also a member of the Wellcome Pain Consortium (Ref. 102645), and the project was supported by a strategic award from the Wellcome (Ref. 102645). S.P.F. is supported by the European Research Council under the European Union's Seventh Framework Programme (Developing Human Connectome Project: FP/2007–2013/ERC Grant Agreement no. 319456). E.P.D. has been supported by the Developing Human Connectome Project (European Research Council Synergy grant FP/2007–2013), and the SSNAP charity, Oxford (Support for the Sick Newborn and their Parents). M.W.W. is funded by the [Wellcome Trust](#) (092753), and is supported by the National Institute for Health Research (NIHR) Oxford Biomedical Research Centre based at Oxford University Hospitals Trust, Oxford University (the views expressed are those of the author(s) and not necessarily those of the NHS, the NIHR or



the Department of Health). Finally, the Wellcome Centre for Integrative Neuroimaging is supported by core funding from the Wellcome Trust (203139/Z/16/Z).

## Appendix A

### A1. Alternative approaches

There are now several methods that characterise resting brain activity in terms of functional modes, both at the group and subject level. The standard pipeline is essentially a two-step process, where the group-level modes are estimated before some form of back-projection is used to extract subject-specific versions of these. Dual regression and related variants thereof, typically combined or integrated with a group-level spatial independent component analysis (ICA) (Beckmann et al., 2005; Calhoun et al., 2001), have been the de facto standard for analyses of the subject variability in spatial and, more recently, temporal features of modes for at least the past decade. Dual regression proceeds by regressing the group-level spatial maps into the data to get a set of time courses—from which subject variability in temporal features may be estimated via any number of functional connectivity metrics—before regressing the time courses back into the data to get subject-specific spatial maps (Beckmann et al., 2009; Calhoun et al., 2001; Erhardt et al., 2011; Nickerson et al., 2017; Zuo et al., 2010b).

This approach has been extended over the years, with several proposed refinements to either the method for identifying group-level modes (Allen et al., 2014; Damoiseaux et al., 2006; Dohmatob et al., 2016; Hjelm et al., 2014; Karahanoğlu and Van De Ville, 2015; Lee et al., 2011; Smith et al., 2012; Varoquaux et al., 2010), or to the way subject-specific information is extracted (Du and Fan, 2013; Hacker et al., 2013; Zöllner et al., 2019).

However, there have been several more extensive departures from the above framework that are more similar in spirit to the hierarchical PFMs model. For example, Abraham et al. (2013); Varoquaux et al. (2011) proposed a more holistic model that finds a set of systems regularised by not only the group-level properties, but also by the consistency of both spatial and temporal information at the subject level. More recently, Li et al. (2017) introduced a model based on non-negative matrix factorisation (NMF) that jointly optimises subject-specific decompositions such that the spatial maps are both sparse and consistent over subjects, though without explicitly leveraging any information about temporal consistency.

As mentioned in the Introduction, these methods all have potential shortcomings in terms of the extent to which typical patterns of variability are learnt from the multiple subject-specific decompositions. These shortcomings are particularly apparent for dual regression type approaches, where the estimation of subject variability is completely post-hoc (and, moreover, the estimated subject variability in spatial organisation only indirectly informs the subject variability in temporal features), but it is also problematic for the more complex models which we have mentioned, for which no explicit parameterisation for the observed variability over subjects is inferred.

More recent methodological work has focused on deriving subject-specific parcellations, both based on a fixed group-level template (Chong et al., 2017; Dhillon et al., 2014; Glasser et al., 2016a; Gordon et al., 2017a; Salehi et al., 2017; Wang et al., 2015), and formulated as a hierarchical model (Kong et al., 2018; Langs et al., 2016; Liu et al., 2012). However, while both mode- and parcel-based approaches have shown promise, our concern is that the subject variability in spatial organisation that has been reported often features reorganisations of a similar scale to our current best estimates of the sizes of distinct functional regions (Van Essen and Dierker, 2007; Van Essen et al., 2012a), and as such, reliable identification at this scale is arguably beyond all but the most sophisticated, multimodal approaches utilising high quality data (Glasser et al., 2016a). Therefore, in this work, we stick to a system-level description and base our method on a decomposition into a set of modes.

Intuitively, a functional system is more protected from the deleterious effects of misalignment than a functional region in two key ways: firstly, functional systems have a greater spatial extent than parcels; secondly, a reorganisation of one region within a larger system can be straightforwardly corrected for if the other regions are relatively stable.

### A2. Preprocessing

The aim of the preprocessing pipeline is to normalise the data such that it has a consistent scale across subjects, and that the properties of the unstructured noise follow the assumptions that are contained in the generative model. The approach is as follows.

- **Voxelwise normalisation.** For each voxel independently, the time course (i.e.  $D_v^{(sp)}$ ) is set to zero mean and unit variance. This ensures that each voxel has a roughly equal contribution to the SVD in the next step.
- **Voxelwise normalisation of the noise subspace.** Each voxel is independently normalised such that the variance of the unstructured noise is unity. This matches the assumption of isotropic noise in the generative model. The unstructured noise subspace is estimated via the SVD. The whole data matrix is decomposed and the  $M$  components with the highest singular values are assumed to represent the structured signal subspace and are removed. The noise subspace is reconstructed from the remaining components, the variance is calculated in each voxel, and the data is renormalised on a voxelwise basis such that the variance becomes unity.
- **Global normalisation of the signal subspace.** There is one final degree of freedom remaining. The generative model assumes isotropic noise, but does not assume a fixed variance. Therefore we can apply a global renormalisation to set the overall variance of the modes we observe. As an approximation, if  $D = PHA$  and we assume independence over modes, then we can say that  $E[D_{vt}^2] = \sum_{m=1}^M E[P_{vm}^2] E[h_m^2] E[A_{mt}^2]$ . In other words, if the maps, amplitudes and time courses have unit variance then the signal variance will be equal to  $M$ . Therefore we use another SVD decomposition and set the overall variance of the assumed signal subspace (i.e. the first  $M$  components) to match the above by applying exactly the same normalisation to each voxel.

### A3. Data reduction

The scale of modern rfMRI studies is now such that even manipulating all the data in its raw form simultaneously is impossible. For approaches that start by inferring group-level descriptions of the data, such as ICA, it is possible to use on-line algorithms that work by passing over the data sequentially (Mensch et al., 2017; Smith et al., 2014), thereby removing the dependence between memory required and the number of subjects under study. However, our approach is explicitly designed to simultaneously extract group- and subject-level features, and as such we need the data from each subject to be available.

To facilitate analyses of large data-sets, we apply subject specific data reductions, but do not collapse these down further to group-level summaries. The approach we take is to approximate each run with a low-rank singular value decomposition (SVD). As our model is defined in terms of both spatial and temporal features, we have to retain both the spatial and temporal singular vectors. However, as the PFM model assumes that subject-specific spatial maps are conserved across all runs for a given subject, we make further savings by only maintaining a single set of spatial singular vectors per subject.

To do this, we calculate the SVD of the matrix formed by temporally concatenating all data from a given subject. This combined data matrix,  $D^{(s)} \in \mathbb{R}^{V \times R_s T}$ , is then represented by  $U^{(s)}$ ,  $S^{(s)}$  and  $V^{(s)}$ . To approximate this with a low rank SVD, we simply only retain the singular vectors associated with the top  $N$  singular values. For example, assuming  $V > R_s T$  and ignoring columns associated with singular values equal



to zero,  $\mathbf{U}^{(s)} \in \mathbb{R}^{V \times R, T}$  is replaced by  $\hat{\mathbf{U}}^{(s)} \in \mathbb{R}^{V \times N}$ . Finally, we can partition the temporal singular vectors, according to the order the individual runs were concatenated, in order to reconstruct the data from each run individually, or in other words,  $\hat{\mathbf{V}}^{(s)} \in \mathbb{R}^{R, T \times N}$  is decomposed into a set of  $\hat{\mathbf{V}}^{(sr)} \in \mathbb{R}^{T \times N}$ . In summary, each data matrix,  $\mathbf{D}^{(sr)}$ , has three approximating matrices, namely  $\hat{\mathbf{U}}^{(s)} \in \mathbb{R}^{V \times N}$ ,  $\hat{\mathbf{S}}^{(s)} \in \mathbb{R}^{N \times N}$  and  $\hat{\mathbf{V}}^{(sr)} \in \mathbb{R}^{T \times N}$ .

The last thing we do is to combine these three matrices into two matrices. This simply saves some computation each time we need to calculate any expectations involving the data. The final form for the approximate data is therefore

$$\begin{aligned} \mathbf{W}^{(s)} &= (\hat{\mathbf{U}}^{(s)}) (\hat{\mathbf{S}}^{(s)})^{\frac{1}{2}} \\ \mathbf{X}^{(sr)} &= (\hat{\mathbf{S}}^{(s)})^{\frac{1}{2}} (\hat{\mathbf{V}}^{(sr)})^T \\ \mathbf{D}^{(sr)} &\approx \mathbf{W}^{(s)} \mathbf{X}^{(sr)} \end{aligned} \quad (20)$$

We can simply substitute this approximate expression for  $\mathbf{D}^{(sr)}$  any time we need access to the data in the inference procedure, and this has the added bonus of being computationally, as well as space, efficient. However, we explicitly calculate, and cache, the overall data variance from the full data, rather than ignoring the contribution from the subspace of discarded singular values<sup>14</sup>. This means that the estimate for the noise precision,  $\psi^{(sr)}$ , will be comparable whether or not we choose to utilise this low-rank approximation, or indeed across different values of  $N$ .

We now have an explicit method for reducing large data-sets to a more manageable size. However, there is one final complication: computationally, calculating the SVD of every  $\mathbf{D}^{(s)}$  actually turns out to be prohibitively expensive in most cases. In order to circumvent this, we utilise the fact that we are explicitly looking for a low-rank approximation and implement an extremely efficient randomised algorithm to directly calculate the truncated SVD. This approach is described in the excellent review by [Halko et al. \(2011\)](#).

#### A4. Degrees of freedom correction

fMRI data has an inherent spatial smoothness—such that there are non-trivial spatial autocorrelations in the noise processes—which is amplified by the spatial smoothing that is a standard pre-processing step for most analyses. As discussed earlier, this is not acknowledged in our specification of our model of the noise process. In essence, this means that the model assumes that there are more independent spatial measurements than actually exist.

Fortunately, as [Groves et al. \(2011\)](#) discuss, there is a simple way to mitigate some of the effects of this within the Bayesian framework. Intuitively, if we have smoothed the data then we should be able to down-sample it without loss of information. At some stage, this would result in the noise becoming genuinely spatially independent again, thereby satisfying the assumptions of the generative model. However, this presents several practical problems, so rather than actually down-sample the data, we simply downweight the spatial information by a factor  $v$ . This represents the proportion of voxels that would be retained if we were to optimally down-sample. This is analogous to fixing that only a random fraction of the data points will be kept, but at each stage averaging over all possible choices of decimated voxels' ([Groves et al., 2011](#)).

While this approach still does not explicitly acknowledge the relationship between noise in nearby voxels, it does counter most of the deleterious effects of this model misspecification, especially when combined with the models for noise in the subject-specific spatial maps and time courses. The main advantage of this approach, compared to a more formal model for smoothness, is that it remains particularly computationally efficient.

<sup>14</sup> More explicitly, we use  $\text{Tr}((\mathbf{D}^{(sr)})^T \mathbf{D}^{(sr)})$  rather than  $\text{Tr}((\mathbf{W}^{(s)} \mathbf{X}^{(sr)})^T \mathbf{W}^{(s)} \mathbf{X}^{(sr)})$  whenever required.

#### A5. Initialisation

With a model of this complexity, it is important that the algorithm is appropriately initialised. By doing so, we can improve reliability and computational stability whilst reducing the computational time required for convergence. Our approach is to compute a consensus group-level set of modes, and use these to initialise the full model.

To do this, we mimic the temporal concatenation approach employed by most existing algorithms and compute a consensus set of spatial singular vectors (this can be done even more efficiently if we have already utilised the data reduction technique described previously ([Calhoun et al., 2001](#))) using another randomised SVD algorithm that streams over the data. These singular vectors are reweighted via an adjusted set of singular values. More specifically, we use the properties of the Marchenko-Pastur distribution to find the noise level that ensures the SNR at the group-level decomposition is similar to the SNR at the subject level. We then run a Bayesian version of spatial ICA—with the spatial priors set to mimic the group-level priors of the full model—to generate the group-level modes. The SNR recalibration ensures we do not get over-splitting of the modes at this stage. We can then propagate this set of group-level modes through the rest of the algorithm, thereby ensuring all parameters are initialised with plausible values.

#### A6. Simulations

The spatial model consists of two levels: parcels and modes. We simulate 100 spatially contiguous parcels within a one-dimensional space comprising 10,000 voxels. We then apply a random diffeomorphic warp to each subject separately as a model for residual misalignments after registration. We then simulate a set of 15 modes consisting of blocks of spatially adjacent parcels. There is variability in the mode weights over subjects, and we introduce overlap such that, on average, each voxel is a member of 1.4 modes.

In the temporal domain we simulate a set of sparse, correlated 'neural' timecourses for each of the two runs per subject. There is variability in the between-mode correlation structure at the run and subject level. These are then convolved with a random draw from the FLOBS basis of haemodynamic response functions ([Woolrich et al., 2004](#)), which introduces variability over subjects and space. This results in 500 timepoints at a TR of 2.0s.

The spatial maps and timecourses are combined via the outer product model, and a nonlinear saturation is applied such that the highest amplitude moments of instantaneous voxelwise activity are reduced. Finally, random noise is added with a degree of spatiotemporal smoothness such that the overall SNR (expressed in terms of variance) is 0.1. For the simulations presented in the Supplementary Material, we also add some structured, subject-specific noise components, again using an outer-product model. These can either be spatially specific or global, and are designed to contribute a similar amount of variance per component as the individual signal modes.

To allow inference algorithms to model the aforementioned artefacts, 18 modes are inferred. After inference, modes are paired to the ground truth based on the similarity between both the spatial maps and timecourses, averaged over runs and subjects. The full set of performance metrics shown in the Supplementary Material is then calculated.

#### A7. Human connectome project data and analyses

For the HCP analyses, all data was from the 1200 Subjects Data Release: [humanconnectome.org/study/hcp-young-adult/document/1200-subjects-data-release](http://humanconnectome.org/study/hcp-young-adult/document/1200-subjects-data-release). We used the 1003 subjects for whom there was full behavioural, structural and rfMRI data (i.e. 4 runs, each of 1200 volumes). All analyses are of the MSMAll and FIX cleaned data (i.e. rfMRI\_REST1\_LR Atlas MSMAll\_hp2000\_clean.dtseries.nii Tr).

ICA-DR results were taken from the Extensively Processed fMRI Data: [humanconnectome.org/study/hcp-young-adult/document/extensively-processed-fmri-data-documentation](https://humanconnectome.org/study/hcp-young-adult/document/extensively-processed-fmri-data-documentation). The amplitudes and netmats were estimated from the time courses released by the HCP: the amplitudes were taken as the standard deviations, while the netmats were the partial correlation matrices. Tikhonov regularisation was used when calculating the inverse of the full correlation matrices, with  $\Gamma = 0.1I$ .

Heritability was estimated via Falconer's formula,  $H_b^2 = 2(r_{mz} - r_{dz})$  (Falconer, 1960). We calculate the correlations,  $r_{mz}$  and  $r_{dz}$ , between the voxelwise spatial map weights. In other words, for each subject and each voxel we extract a length  $M$  vector of weights:  $\mathbf{P}_v^{(s)}$  using the PFM notation, and compute the correlation between these for every pair of subjects.

For the CCA analyses, we used the full set of restricted information released by the HCP. We first removed all variables relating to study completion or quality control. The structural variables were all the remaining variables in the FreeSurfer category; all others were taken as behavioural. To preprocess the behavioural variables, we first removed any variables that were either more than 20% NaN, or those for which more than 95% of subjects had exactly the same entries. We then imputed any missing values using the SoftImpute method (Mazumder et al., 2010) as implemented in the fancyimpute Python package ([github.com/iskandr/fancyimpute](https://github.com/iskandr/fancyimpute)). The following were regressed out as confounds (one-hot encoded where necessary) in all subsequent analyses: Release, Acquisition, fMRI\_3T\_ReconVrs, rfMRI\_motion, Age, Gender, Race, Ethnicity, Handedness, Height, Weight, BMI, BPSystolic, BPDiastolic, Hematocrit\_1, Hematocrit\_2, FS\_IntraCranial\_Vol, FS\_BrainSeg\_Vol. A detailed description of all variables can be found at [wiki.humanconnectome.org/display/PublicData/HCP+Data+Dictionary+Public+Updated+for+the+1200+Subject+Release](https://wiki.humanconnectome.org/display/PublicData/HCP+Data+Dictionary+Public+Updated+for+the+1200+Subject+Release).

For the CCA, all groups of variables were normalised and then reduced to their top 25 components via the SVD, before a CCA was run on every pair of variable groups. The RV coefficient was then calculated between the top 10 paired components from each CCA.

#### A8. Active-state data and analyses

Data was acquired from fifteen subjects, but for these analyses we excluded Subject 07 due to potential artefacts in several of their scans. Preprocessing was as previously published (i.e. brain extraction, B0 unwarping, high-pass temporal filtering, motion correction, and FIX cleaning) (Kieliba et al., 2019). However, we did not apply mean-based intensity normalisation or low-pass filter the data. Finally, the preprocessed functional scans were then registered to MNI space and spatially smoothed (2mm FWHM).

As with the HCP data, the ICA-DR amplitudes and netmats were estimated from the time courses: the amplitudes were taken as the standard deviations, while the netmats were the partial correlation matrices. Tikhonov regularisation was used when calculating the inverse of the full correlation matrices, with  $\Gamma = 0.1I$ . The SVM was from scikit-learn (`sklearn.svm.SVC`), and as this is a relatively small dataset parameters were left at their defaults (Varoquaux et al., 2017).

#### Supplementary material

Supplementary material associated with this article can be found, in the online version, at [10.1016/j.neuroimage.2020.117226](https://doi.org/10.1016/j.neuroimage.2020.117226)

#### References

Abraham, A., Dohmatob, E., Thirion, B., Samaras, D., Varoquaux, G., 2013. Extracting Brain Regions from Rest fMRI with Total-variation Constrained Dictionary Learning. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (Eds.), Medical Image Computing and Computer-Assisted Intervention—MICCAI 2013. In: Lecture Notes in Computer Science, 8150. Springer Berlin Heidelberg, pp. 607–615. doi:[10.1007/978-3-642-40763-5\\_75](https://doi.org/10.1007/978-3-642-40763-5_75).

Abraham, A., Milham, M.P., Di Martino, A., Craddock, R.C., Samaras, D., Thirion, B., Varoquaux, G., 2017. Deriving reproducible biomarkers from multi-site resting-state data: an autism-based example. *Neuroimage* 147, 736–745. doi:[10.1016/j.neuroimage.2016.10.045](https://doi.org/10.1016/j.neuroimage.2016.10.045).

Allen, E.A., Erhardt, E.B., Damaraju, E., Gruner, W., Segall, J.M., Silva, R.F., Havlicek, M., Rachakonda, S., Fries, J., Kalyanam, R., Michael, A.M., Caprihan, A., Turner, J.A., Eichele, T., Adelsheim, S., Bryan, A.D., Bustillo, J., Clark, V.P., Feldstein Ewing, S.W., Filbey, F., Ford, C.C., Hutchison, K., Jung, R.E., Kiehl, K.A., Kodituwakku, P., Komesu, Y.M., Mayer, A.R., Pearlson, G.D., Phillips, J.P., Sadek, J.R., Stevens, M., Teuscher, U., Thoma, R.J., Calhoun, V.D., 2011. A baseline for the multivariate comparison of resting state networks. *Front. Syst. Neurosci.* 5 (2). doi:[10.3389/fn-sys.2011.00002](https://doi.org/10.3389/fn-sys.2011.00002).

Allen, E.A., Erhardt, E.B., Wei, Y., Eichele, T., Calhoun, V.D., 2012. Capturing inter-subject variability with group independent component analysis of fmri data: a simulation study. *Neuroimage* 59 (4), 4141–4159. doi:[10.1016/j.neuroimage.2011.10.010](https://doi.org/10.1016/j.neuroimage.2011.10.010).

Allen, G.I., Grosenick, L., Taylor, J., 2014. A generalized least-square matrix decomposition. *J. Am. Stat. Assoc.* 109 (505), 145–159. doi:[10.1080/01621459.2013.852978](https://doi.org/10.1080/01621459.2013.852978).

Amiez, C., Petrides, M., 2014. Neuroimaging evidence of the anatomo-functional organization of the human cingulate motor areas. *Cerebral Cortex* 24 (3), 563–578. doi:[10.1093/cercor/bhs329](https://doi.org/10.1093/cercor/bhs329).

Amunts, K., Schleicher, A., Zilles, K., 2007. Cytoarchitecture of the cerebral cortex more than localization. *Neuroimage* 37 (4), 1061–1065. doi:[10.1016/j.neuroimage.2007.02.037](https://doi.org/10.1016/j.neuroimage.2007.02.037).

Andrews, T.J., Halpern, S.D., Purves, D., 1997. Correlated size variations in human visual cortex, lateral geniculate nucleus, and optic tract. *The Journal of Neuroscience* 17 (8), 2859–2868.

Anzellotti, S., Coutanche, M.N., 2018. Beyond functional connectivity: investigating networks of multivariate representations. *Trends Cogn. Sci. (Regul. Ed.)* 22 (3), 258–269. doi:[10.1016/j.tics.2017.12.002](https://doi.org/10.1016/j.tics.2017.12.002).

Arichi, T., Fagiolo, G., Varela, M., Melendez-Calderon, A., Allievi, A., Merchant, N., Tumor, N., Counsell, S.J., Burdet, E., Beckmann, C.F., Edwards, A.D., 2012. Development of BOLD signal hemodynamic responses in the human brain. *Neuroimage* 63 (2), 663–673. doi:[10.1016/j.neuroimage.2012.06.054](https://doi.org/10.1016/j.neuroimage.2012.06.054).

Attias, H., 2000. A Variational Bayesian Framework for Graphical Models. In: Solla, S.A., Leen, T.K., Müller, K. (Eds.), *Advances in Neural Information Processing Systems 12*. MIT Press, pp. 209–215.

Bamberg, F., Kauczor, H.-U., Weckbach, S., Schlett, C.L., Forsting, M., Ladd, S.C., Greiser, K.H., Weber, M.-A., Schulz-Menger, J., Niendorf, T., Pischon, T., Caspers, S., Amunts, K., Berger, K., Blow, R., Hosten, N., Hegenscheid, K., Krncke, T., Linseisen, J., Gnther, M., Hirsch, J.G., Khn, A., Hendel, T., Wichmann, H.-E., Schmidt, B., Jckel, K.-H., Hoffmann, W., Kaaks, R., Reiser, M.F., Vlzke, H., For the German National Cohort MRI Study Investigators, 2015. Whole-body MR imaging in the german national cohort: Rationale, design, and technical background. *Radiology* 277 (1), 206–220. doi:[10.1148/radiol.2015142272](https://doi.org/10.1148/radiol.2015142272). PMID: 25989618.

Bartley, A.J., Jones, D.W., Weinberger, D.R., 1997. Genetic variability of human brain size and cortical gyral patterns. *Brain* 120 (2), 257–269. doi:[10.1093/brain/120.2.257](https://doi.org/10.1093/brain/120.2.257).

Beckmann, C.F., DeLuca, M., Devlin, J.T., Smith, S.M., 2005. Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions of the Royal Society B: Biological Sciences* 360 (1457), 1001–1013. doi:[10.1098/rstb.2005.1634](https://doi.org/10.1098/rstb.2005.1634).

Beckmann, C.F., Mackay, C., Filippini, N., Smith, S.M., 2009. Group comparison of resting-state fMRI data using multi-subject ICA and dual regression. *Organization for Human Brain Mapping*.

Beckmann, C.F., Smith, S.M., 2004. Probabilistic independent component analysis for functional magnetic resonance imaging. *Medical Imaging, IEEE Transactions on* 23 (2), 137–152. doi:[10.1109/TMI.2003.822821](https://doi.org/10.1109/TMI.2003.822821).

Bijsterbosch, J.D., Beckmann, C.F., Woolrich, M.W., Smith, S.M., Harrison, S.J., 2019. The relationship between spatial configuration and functional connectivity of brain regions revisited. *Elife* 8, e44890. doi:[10.7554/eLife.44890](https://doi.org/10.7554/eLife.44890).

Bijsterbosch, J.D., Harrison, S.J., Duff, E.P., Alfaro-Almagro, F., Woolrich, M.W., Smith, S.M., 2017. Investigations into within- and between-subject resting-state amplitude variations. *Neuroimage* 159, 57–69. doi:[10.1016/j.neuroimage.2017.07.014](https://doi.org/10.1016/j.neuroimage.2017.07.014).

Bijsterbosch, J.D., Woolrich, M.W., Glasser, M.F., Robinson, E.C., Beckmann, C.F., Van Essen, D.C., Harrison, S.J., Smith, S.M., 2018. The relationship between spatial configuration and functional connectivity of brain regions. *Elife* 7, e32992. doi:[10.7554/eLife.32992](https://doi.org/10.7554/eLife.32992).

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., 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. *Proceedings of the National Academy of Sciences* 107 (10), 4734–4739. doi:[10.1073/pnas.0911855107](https://doi.org/10.1073/pnas.0911855107).

Blei, D.M., Kucukelbir, A., McAuliffe, J.D., 2017. Variational inference: a review for statisticians. *J. Am. Stat. Assoc.* 112 (518), 859–877. doi:[10.1080/01621459.2017.1285773](https://doi.org/10.1080/01621459.2017.1285773).

Braga, R.M., Buckner, R.L., 2017. Parallel interdigitated distributed networks within the individual estimated by intrinsic functional connectivity. *Neuron* 95 (2), 457–471.e5. doi:[10.1016/j.neuron.2017.06.038](https://doi.org/10.1016/j.neuron.2017.06.038).

Breteler, M.M.B., Stöcker, T., Pracht, E., Brenner, D., Stirnberg, R., 2014. MRI in the Rhineland study: A novel protocol for population neuroimaging. *Alzheimer's & De-*

- mentia 10 (4), 92. doi:10.1016/j.jalz.2014.05.172. Alzheimer's Association International Conference 2014.
- Brett, M., Johnsrude, I.S., Owen, A.M., 2002. The problem of functional localization in the human brain. *Nat. Rev. Neurosci.* 3 (3), 243–249. doi:10.1038/nrn756.
- Bright, M.G., Murphy, K., 2015. Is fMRI “noise” really noise? resting state nuisance regressors remove variance with network structure. *Neuroimage* 114, 158–169. doi:10.1016/j.neuroimage.2015.03.070.
- Buckner, R.L., Andrews-Hanna, J.R., Schacter, D.L., 2008. The brain's default network: anatomy, function, and relevance to disease. *Ann. N. Y. Acad. Sci.* 1124 (1), 1–38. doi:10.1196/annals.1440.011.
- Buckner, R.L., Krienen, F.M., Yeo, B.T.T., 2013. Opportunities and limitations of intrinsic functional connectivity MRI. *Nat. Neurosci.* 16 (7), 832–837. doi:10.1038/nn.3423.
- Calhoun, V., Miller, R., Pearson, G., Adali, T., 2014. The chronnectome: time-varying connectivity networks as the next frontier in fMRI data discovery. *Neuron* 84 (2), 262–274. doi:10.1016/j.neuron.2014.10.015.
- Calhoun, V.D., Adali, T., Pearlson, G.D., Pekar, J.J., 2001. A method for making group inferences from functional MRI data using independent component analysis. *Hum. Brain Mapp.* 14 (3), 140–151. doi:10.1002/hbm.1048.
- Calhoun, V.D., Potluru, V.K., Phlypo, R., Silva, R.F., Pearlmutter, B.A., Caprihan, A., Plis, S.M., Adali, T., 2013. Independent component analysis for brain fMRI does indeed select for maximal independence. *PLoS ONE* 8 (8), e73309. doi:10.1371/journal.pone.0073309.
- Chong, M., Bhushan, C., Joshi, A.A., Choi, S., Halder, J.P., Shattuck, D.W., Spreng, R.N., Leahy, R.M., 2017. Individual parcellation of resting fMRI with a group functional connectivity prior. *Neuroimage* 156, 87–100. doi:10.1016/j.neuroimage.2017.04.054.
- Coalson, T.S., Van Essen, D.C., Glasser, M.F., 2018. The impact of traditional neuroimaging methods on the spatial localization of cortical areas. *Proceedings of the National Academy of Sciences* 115 (27), E6356–E6365. doi:10.1073/pnas.1801582115.
- Colclough, G.L., Smith, S.M., Nichols, T.E., Winkler, A.M., Sotiropoulos, S.N., Glasser, M.F., Van Essen, D.C., Woolrich, M.W., 2017. The heritability of multi-modal connectivity in human brain activity. *Elife* 6, e20178. doi:10.7554/eLife.20178.
- Colclough, G.L., Woolrich, M.W., Harrison, S.J., Rojas López, P.A., Valdes-Sosa, P.A., Smith, S.M., 2018. Multi-subject hierarchical inverse covariance modelling improves estimation of functional brain networks. *Neuroimage* 178, 370–384. doi:10.1016/j.neuroimage.2018.04.077.
- Conroy, B.R., Singer, B.D., Guntupalli, J.S., Ramadge, P.J., Haxby, J.V., 2013. Inter-subject alignment of human cortical anatomy using functional connectivity. *Neuroimage* 81, 400–411. doi:10.1016/j.neuroimage.2013.05.009.
- Couvy-Duchesne, B., Blokland, G.A.M., Hickie, I.B., Thompson, P.M., Martin, N.G., de Zubicaray, G.I., McMahon, K.L., Wright, M.J., 2014. Heritability of head motion during resting state functional MRI in 462 healthy twins. *Neuroimage* 102, Part 2, 424–434. doi:10.1016/j.neuroimage.2014.08.010.
- Dadi, K., Rahim, M., Abraham, A., Chyzykh, D., Milham, M., Thirion, B., Varoquaux, G., 2019. Benchmarking functional connectome-based predictive models for resting-state fMRI. *Neuroimage* 192, 115–134. doi:10.1016/j.neuroimage.2019.02.062.
- 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. *Proceedings of the National Academy of Sciences* 103 (37), 13848–13853. doi:10.1073/pnas.0601417103.
- Devlin, J.T., Poldrack, R.A., 2007. In praise of tedious anatomy. *Neuroimage* 37 (4), 1033–1041. doi:10.1016/j.neuroimage.2006.09.055.
- Dhillon, P.S., Wolk, D.A., Das, S.R., Ungar, L.H., Gee, J.C., Avants, B.B., 2014. Subject-specific functional parcellation via prior based eigenanatomy. *Neuroimage* 99, 14–27. doi:10.1016/j.neuroimage.2014.05.026.
- Dohmatob, E., Mensch, A., Varoquaux, G., Thirion, B., 2016. Learning Brain Regions via Large-scale Online Structured Sparse Dictionary Learning. In: Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R. (Eds.), *Advances in Neural Information Processing Systems 29*. Curran Associates, Inc., pp. 4610–4618.
- Douaud, G., Groves, A.R., Tamnes, C.K., Westlye, L.T., Duff, E.P., Engvig, A., Walhovd, K.B., James, A., Gass, A., Monsch, A.U., Matthews, P.M., Fjell, A.M., Smith, S.M., Johansen-Berg, H., 2014. A common brain network links development, aging, and vulnerability to disease. *Proceedings of the National Academy of Sciences* 111 (49), 17648–17653. doi:10.1073/pnas.1410378111.
- Dougherty, R.F., Koch, V.M., Brewer, A.A., Fischer, B., Modersitzki, J., Wandell, B.A., 2003. Visual field representations and locations of visual areas v1/2/3 in human visual cortex. *J. Vis.* 3 (10), 586–598. doi:10.1167/3.10.1.
- Du, Y., Fan, Y., 2013. Group information guided ICA for fMRI data analysis. *Neuroimage* 69, 157–197. doi:10.1016/j.neuroimage.2012.11.008.
- Dubois, J., Adolphs, R., 2016. Building a science of individual differences from fMRI. *Trends Cogn. Sci. (Regul. Ed.)* 20 (6), 425–443. doi:10.1016/j.tics.2016.03.014.
- Duff, E.P., Johnston, L.A., Xiong, J., Fox, P.T., Mareels, I., Egan, G.F., 2008. The power of spectral density analysis for mapping endogenous BOLD signal fluctuations. *Hum. Brain Mapp* 29 (7), 778–790. doi:10.1002/hbm.20601.
- Duff, E.P., Makin, T., Cottaar, M., Smith, S.M., Woolrich, M.W., 2018. Disambiguating brain functional connectivity. *Neuroimage* 173, 540–550. doi:10.1016/j.neuroimage.2018.01.053.
- Eickhoff, S., Nichols, T.E., Van Horn, J.D., Turner, J.A., 2016. Sharing the wealth: Neuroimaging data repositories. *NeuroImage* 124 (Part B), 1065–1068. doi:10.1016/j.neuroimage.2015.10.079. Sharing the wealth: Brain Imaging Repositories in 2015.
- Eklund, A., Nichols, T.E., Knutsson, H., 2016. Cluster failure: why fMRI inferences for spatial extent have inflated false-positive rates. *Proceedings of the National Academy of Sciences* 113 (28), 7900–7905. doi:10.1073/pnas.1602413113.
- Elliott, L.T., Sharp, K., Alfaro-Almagro, F., Shi, S., Miller, K.L., Douaud, G., Marchini, J., Smith, S.M., 2018. Genome-wide association studies of brain imaging phenotypes in UK biobank. *Nature* 562 (7726), 210–216. doi:10.1038/s41586-018-0571-7.
- Erhardt, E.B., Rachakonda, S., Bedrick, E.J., Allen, E.A., Adali, T., Calhoun, V.D., 2011. Comparison of multi-subject ICA methods for analysis of fMRI data. *Hum. Brain Mapp.* 32 (12), 2075–2095. doi:10.1002/hbm.21170.
- Esposito, F., Scarabino, T., Hyvärinen, A., Himberg, J., Formisano, E., Comani, S., Tedeschi, G., Goebel, R., Seifritz, E., Di Salle, F., 2005. Independent component analysis of fMRI group studies by self-organizing clustering. *Neuroimage* 25 (1), 193–205. doi:10.1016/j.neuroimage.2004.10.042.
- Fagerland, M.W., Lydersen, S., Laake, P., 2013. The McNemar test for binary matched-pairs data: mid-p and asymptotic are better than exact conditional. *BMC Med. Res. Methodol.* 13, 91. doi:10.1186/1471-2288-13-91.
- Falconer, D.S., 1960. *Introduction to quantitative genetics*. Oliver & Boyd, Edinburgh/London.
- Finn, E.S., Shen, X., Scheinost, D., Rosenberg, M.D., Huang, J., Chun, M.M., Papademetris, X., Constable, R.T., 2015. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nat. Neurosci.* 18 (11), 1664–1671. doi:10.1038/nn.4135.
- Friston, K.J., 2011. Functional and effective connectivity: a review. *Brain Connect.* 1 (1), 13–36. doi:10.1089/brain.2011.0008.
- , 2007. In: Friston, K.J., Ashburner, J., Kiebel, S., Nichols, T., Penny, W. (Eds.), *Statistical parametric mapping: The analysis of functional brain images*. Academic Press.
- Friston, K.J., Frith, C.D., Liddle, P.F., Frackowiak, R.S.J., 1993. Functional connectivity: The principal-component analysis of large (PET) data sets. *Journal of Cerebral Blood Flow & Metabolism* 13 (1), 5–14. doi:10.1038/jcbfm.1993.4. PMID: 8417010.
- Geerligs, L., Cam-CAN, Henson, R.N., 2016. Functional connectivity and structural covariance between regions of interest can be measured more accurately using multivariate distance correlation. *Neuroimage* 135, 16–31. doi:10.1016/j.neuroimage.2016.04.047.
- Gelman, A., 2006. Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian Analysis* 1 (3), 515–534. doi:10.1214/06-BA117A.
- George, E.I., McCulloch, R.E., 1993. Variable selection via Gibbs sampling. *J. Am. Stat. Assoc.* 88 (423), 881–889. doi:10.1080/01621459.1993.10476353.
- Glasser, M.F., Coalson, T.S., Bijsterbosch, J.D., Harrison, S.J., Harms, M.P., Anticevic, A., Van Essen, D.C., Smith, S.M., 2018. Using temporal ICA to selectively remove global noise while preserving global signal in functional MRI data. *Neuroimage* 181, 692–717. doi:10.1016/j.neuroimage.2018.04.076.
- Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., Smith, S.M., Van Essen, D.C., 2016. A multi-modal parcellation of human cerebral cortex. *Nature* 536, 171–178. doi:10.1038/nature18933.
- Glasser, M.F., Smith, S.M., Marcus, D.S., Andersson, J.L.R., 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. *Nat. Neurosci.* 19 (9), 1175–1187. doi:10.1038/nn.4361.
- 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. doi:10.1016/j.neuroimage.2013.04.127. Mapping the Connectome.
- Glasser, M.F., Van Essen, D.C., 2011. Mapping human cortical areas in vivo based on myelin content as revealed by T1- and T2-weighted MRI. *The Journal of Neuroscience* 31 (32), 11597–11616. doi:10.1523/JNEUROSCI.2180-11.2011.
- Gonzalez-Castillo, J., Saad, Z.S., Handwerker, D.A., Inati, S.J., Brenowitz, N., Bandettini, P.A., 2012. Whole-brain, time-locked activation with simple tasks revealed using massive averaging and model-free analysis. *Proceedings of the National Academy of Sciences* 109 (14), 5487–5492. doi:10.1073/pnas.1121049109.
- Gordon, E.M., Laumann, T.O., Adeyemo, B., Gilmore, A.W., Nelson, S.M., Dosenbach, N.U.F., Petersen, S.E., 2016. Individual-specific features of brain systems identified with resting state functional correlations. *Neuroimage* 146, 918–939. doi:10.1016/j.neuroimage.2016.08.032.
- Gordon, E.M., Laumann, T.O., Adeyemo, B., Petersen, S.E., 2017. Individual variability of the system-level organization of the human brain. *Cerebral Cortex* 27 (1), 386–399. doi:10.1093/cercor/bhv239.
- Gordon, E.M., Laumann, T.O., Gilmore, A.W., Newbold, D.J., Greene, D.J., Berg, J.J., Ortega, M., Hoyt-Drazen, C., Grattton, C., Sun, H., Hampton, J.M., Coalson, R.S., Nguyen, A.L., McDermott, K.B., Shimony, J.S., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., Nelson, S.M., Dosenbach, N.U.F., 2017. Precision functional mapping of individual human brains. *Neuron* 95 (4), 791–807. doi:10.1016/j.neuron.2017.07.011.
- Gorgolewski, K.J., Esteban, O., Schafer, G., Wandell, B., Poldrack, R.A., 2017. *Openneuro - a free online platform for sharing and analysis of neuroimaging data*. *Organization for Human Brain Mapping*.
- Grattton, C., Laumann, T.O., Nielsen, A.N., Greene, D.J., Gordon, E.M., Gilmore, A.W., Nelson, S.M., Coalson, R.S., Snyder, A.Z., Schlaggar, B.L., Dosenbach, N.U.F., Petersen, S.E., 2018. Functional brain networks are dominated by stable group and individual factors, not cognitive or daily variation. *Neuron* 98 (2), 439–452. doi:10.1016/j.neuron.2018.03.035.
- Greicius, M.D., Krasnow, B., Reiss, A.L., Menon, V., 2003. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences* 100 (1), 253–258. doi:10.1073/pnas.0135058100.
- Groves, A.R., Beckmann, C.F., Smith, S.M., Woolrich, M.W., 2011. Linked independent component analysis for multimodal data fusion. *Neuroimage* 54 (3), 2198–2217. doi:10.1016/j.neuroimage.2010.09.073.
- Guntupalli, J.S., Feilong, M., Haxby, J.V., 2018. A computational model of shared fine-scale structure in the human connectome. *PLoS Comput. Biol.* 14 (4), 1–26. doi:10.1371/journal.pcbi.1006120.



- Guntupalli, J.S., Hanke, M., Halchenko, Y.O., Connolly, A.C., Ramadge, P.J., Haxby, J.V., 2016. A model of representational spaces in human cortex. *Cerebral Cortex* 26 (6), 2919–2934. doi:10.1093/cercor/bhw068. /oup/backfile/Content\_public/Journal/cercor/26/6/10.1093\_cercor\_bhw068/3/bhw068.pdf.
- Haak, K.V., Marquand, A.F., Beckmann, C.F., 2018. Connectopic mapping with resting-state fMRI. *NeuroImage* 170, 83–94. doi:10.1016/j.neuroimage.2017.06.075. Segmenting the Brain.
- Hacker, C.D., Laumann, T.O., Szrama, N.P., Baldassarre, A., Snyder, A.Z., Leuthardt, E.C., Corbetta, M., 2013. Resting state network estimation in individual subjects. *NeuroImage* 82, 616–633. doi:10.1016/j.neuroimage.2013.05.108.
- Haldane, J.B.S., 1932. A note on inverse probability. *Math. Proc. Cambridge Philos. Soc.* 28 (1), 55–61. doi:10.1017/S0305004100010495.
- Halko, N., Martinsson, P.G., Tropp, J.A., 2011. Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.* 53 (2), 217–288. doi:10.1137/090771806.
- Harrison, S.J., Woolrich, M.W., Robinson, E.C., Glasser, M.F., Beckmann, C.F., Jenkinson, M., Smith, S.M., 2015. Large-scale probabilistic functional modes from resting state fmri. *NeuroImage* 109, 217–231. doi:10.1016/j.neuroimage.2015.01.013.
- Hasan, A., McIntosh, A.M., Droese, U.-A., Schneider-Axmann, T., Lawrie, S.M., Moorhead, T.W., Tepest, R., Maier, W., Falkai, P., Wobrock, T., 2011. Prefrontal cortex gyrification index in twins: an MRI study. *Eur. Arch. Psychiatry Clin. Neurosci.* 261 (7), 459–465. doi:10.1007/s00406-011-0198-2.
- Hjelm, R.D., Calhoun, V.D., Salakhutdinov, R., Allen, E.A., Adali, T., Plis, S.M., 2014. Restricted boltzmann machines for neuroimaging: an application in identifying intrinsic networks. *NeuroImage* 96, 245–260. doi:10.1016/j.neuroimage.2014.03.048.
- Hodgson, K., Poldrack, R.A., Curran, J.E., Knowles, E.E., Mathias, S., Göring, H.H.H., Yao, N., Olvera, R.L., Fox, P.T., Almasy, L., Duggirala, R., Barch, D.M., Blangero, J., Glahn, D.C., 2017. Shared genetic factors influence head motion during MRI and body mass index. *Cerebral Cortex* 27 (12), 5539–5546. doi:10.1093/cercor/bhw321.
- Horien, C., Shen, X., Scheinost, D., Constable, R.T., 2019. The individual functional connectome is unique and stable over months to years. *NeuroImage* 189, 676–687. doi:10.1016/j.neuroimage.2019.02.002.
- Hotelling, H., 1936. Relations between two sets of variates. *Biometrika* 28 (3–4), 321–377. doi:10.1093/biomet/28.3-4.321.
- Hutchison, R.M., Womelsdorf, T., Allen, E.A., Bandettini, P.A., Calhoun, V.D., Corbetta, M., Penna, S.D., Duyn, J.H., Glover, G.H., Gonzalez-Castillo, J., Handwerker, D.A., Keilholz, S., Kiviniemi, V., Leopold, D.A., de Pasquale, F., Sporns, O., Walter, M., Chang, C., 2013. Dynamic functional connectivity: Promise, issues, and interpretations. *NeuroImage* 80, 360–378. doi:10.1016/j.neuroimage.2013.05.079. Mapping the Connectome.
- Huth, A.G., de Heer, W.A., Griffiths, T.L., Theunissen, F.E., Gallant, J.L., 2016. Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature* 532 (7600), 453–458. doi:10.1038/nature17637.
- Insel, T.R., Cuthbert, B.N., 2015. Brain disorders? precisely. *Science* 348 (6234), 499–500. doi:10.1126/science.aab2358.
- Ishwaran, H., Rao, J.S., 2005. Spike and slab variable selection: frequentist and bayesian strategies. *The Annals of Statistics* 33 (2), 730–773.
- Jenkinson, M., Beckmann, C.F., Behrens, T.E.J., Woolrich, M.W., Smith, S.M., 2012. FSL. *NeuroImage* 62 (2), 782–790. doi:10.1016/j.neuroimage.2011.09.015. 20 YEARS OF fMRI 20 YEARS OF fMRI.
- Johansen-Berg, H., Behrens, T.E.J., Robson, M.D., Drobnyak, I., Rushworth, M.F.S., Brady, J.M., Smith, S.M., Higham, D.J., Matthews, P.M., 2004. Changes in connectivity profiles define functionally distinct regions in human medial frontal cortex. *Proceedings of the National Academy of Sciences* 101 (36), 13335–13340. doi:10.1073/pnas.0403743101.
- 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). doi:10.1038/ncomms8751.
- Kennedy, D.N., Haselgrove, C., Riehl, J., Preuss, N., Buccigrossi, R., 2016. The NITRC image repository. *NeuroImage* 124 (Part B), 1069–1073. doi:10.1016/j.neuroimage.2015.05.074. Sharing the wealth: Brain Imaging Repositories in 2015.
- Kieliba, P., Madugula, S., Filippini, N., Duff, E.P., Makin, T.R., 2019. Large-scale intrinsic connectivity is consistent across varying task demands. *PLoS ONE* 14 (4), 1–21. doi:10.1371/journal.pone.0213861.
- Kiviniemi, V., Starck, T., Remes, J., Long, X., Nikkinen, J., Haapea, M., Veijola, J., Moilanen, L., Isohanni, M., Zang, Y.-F., Tervonen, O., 2009. Functional segmentation of the brain cortex using high model order group PICA. *Hum. Brain Mapp.* 30 (12), 3865–3886. doi:10.1002/hbm.20813.
- Kong, R., Li, J., Orban, C., Sabuncu, M.R., Liu, H., Schaefer, A., Sun, N., Zuo, X.-N., Holmes, A.J., Eickhoff, S.B., Yeo, B.T.T., 2018. Spatial topography of individual-specific cortical networks predicts human cognition, personality, and emotion. *Cerebral Cortex* doi:10.1093/cercor/bhy123. bhy123 /oup/backfile/content\_public/journal/cercor/pap/10.1093\_cercor\_bhy123/1/bhy123.pdf.
- Kriegeskorte, N., Bodurka, J., Bandettini, P., 2008. Artfactual time-course correlations in echo-planar fmri with implications for studies of brain function. *Int. J. Imaging Syst. Technol.* 18 (5–6), 345–349. doi:10.1002/ima.20166.
- Krienen, F.M., Yeo, B.T.T., Buckner, R.L., 2014. Reconfigurable task-dependent functional coupling modes cluster around a core functional architecture. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 369 (1653), 20130526. doi:10.1098/rstb.2013.0526.
- Kruschke, J., 2014. *Doing bayesian data analysis: A Tutorial with r, JAGS, and stan*, 2nd Academic Press.
- Langs, G., Tie, Y., Rigolo, L., Golby, A., Golland, P., 2010. *Functional Geometry Alignment and Localization of Brain Areas*. In: Lafferty, J.D., Williams, C.K.I., Shawe-Taylor, J., Zemel, R.S., Culotta, A. (Eds.), *Advances in Neural Information Processing Systems 23*. Curran Associates, Inc., pp. 1225–1233.
- Langs, G., Wang, D., Golland, P., Mueller, S., Pan, R., Sabuncu, M.R., Sun, W., Li, K., Liu, H., 2016. Identifying shared brain networks in individuals by decoupling functional and anatomical variability. *Cerebral Cortex* 26 (10), 4004–4014. doi:10.1093/cercor/bhw189. /oup/backfile/Content\_public/Journal/cercor/26/10/10.1093\_cercor\_bhw189/2/bhw189.pdf.
- Laumann, T.O., Gordon, E.M., Adeyemo, B., Snyder, A.Z., Joo, S.J., Chen, M.-Y., Gilmore, A.W., McDermett, K.B., Nelson, S.M., Dosenbach, N.U.F., Schlaggar, B.L., Mumford, J.A., Poldrack, R.A., Petersen, S.E., 2015. Functional system and areal organization of a highly sampled individual human brain. *Neuron* 87 (3), 657–670. doi:10.1016/j.neuron.2015.06.037.
- Laumann, T.O., Snyder, A.Z., Mitra, A., Gordon, E.M., Gratton, C., Adeyemo, B., Gilmore, A.W., Nelson, S.M., Berg, J.J., Greene, D.J., McCarthy, J.E., Tagliazucchi, E., Laufs, H., Schlaggar, B.L., Dosenbach, N.U.F., Petersen, S.E., 2017. On the stability of BOLD fmri correlations. *Cerebral Cortex* 27 (10), 4719–4732. doi:10.1093/cercor/bhw265.
- Lee, J.-H., Hashimoto, R., Wible, C.G., Yoo, S.-S., 2011. Investigation of spectrally coherent resting-state networks using non-negative matrix factorization for functional MRI data. *Int. J. Imaging Syst. Technol.* 21 (2), 211–222. doi:10.1002/ima.20276.
- Li, H., Satterthwaite, T.D., Fan, Y., 2017. Large-scale sparse functional networks from resting state fmri. *NeuroImage* 156, 1–13. doi:10.1016/j.neuroimage.2017.05.004.
- Liu, W., Awate, S.P., Fletcher, P.T., 2012. Group Analysis of Resting-state Fmri by Hierarchical Markov Random Fields. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (Eds.), *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2012*. In: *Lecture Notes in Computer Science*, 7512. Springer Berlin Heidelberg, pp. 189–196. doi:10.1007/978-3-642-33454-2\_24.
- Llera, A., Wolfers, T., Mulders, P., Beckmann, C.F., 2019. Inter-individual differences in human brain structure and morphology link to variation in demographics and behavior. *Elife* 8, e44443. doi:10.7554/eLife.44443.
- MacKay, D.J.C., 2003. *Information theory, inference and learning algorithms*. Cambridge University Press.
- Marcus, D.S., Harms, M.P., Snyder, A.Z., Jenkinson, M., Wilson, J.A., Glasser, M.F., Barch, D.M., Archie, K.A., Burgess, G.C., Ramaratnam, M., Hodge, M., Horton, W., Herrick, R., Olsen, T., McKay, M., House, M., Hileman, M., Reid, E., Harwell, J., Coalson, T., Schindler, J., Elam, J.S., Curtiss, S.W., Van Essen, D.C., 2013. Human Connectome Project informatics: Quality control, database services, and data visualization. *NeuroImage* 80, 202–219. doi:10.1016/j.neuroimage.2013.05.077. Mapping the Connectome.
- Marrelec, G., Krainik, A., Duffau, H., Pélégriani-Issac, M., Lehericy, S., Doyon, J., Benali, H., 2006. Partial correlation for functional brain interactivity investigation in functional MRI. *NeuroImage* 32 (1), 228–237. doi:10.1016/j.neuroimage.2005.12.057.
- Mazumder, R., Hastie, T., Tibshirani, R., 2010. Spectral regularization algorithms for learning large incomplete matrices. *Journal of Machine Learning Research* 11, 2287–2322.
- McDaniel, M.A., 2005. Big-brained people are smarter: a meta-analysis of the relationship between in vivo brain volume and intelligence. *Intelligence* 33 (4), 337–346. doi:10.1016/j.intell.2004.11.005.
- Mennes, M., Biswal, B.B., Castellanos, F.X., Milham, M.P., 2013. Making data sharing work: the FCP/INDI experience. *NeuroImage* 82, 683–691. doi:10.1016/j.neuroimage.2012.10.064.
- Mensch, A., Mairal, J., Thirion, B., Varoquaux, G., 2017. *Stochastic subsampling for factorizing huge matrices*. ArXiv e-prints.
- Miller, K.L., Alfaro-Almagro, F., Bangerter, N.K., Thomas, D.L., Yacoub, E., Xu, J., Bartsch, A.J., Jbabdi, S., Sotiropoulos, S.N., Andersson, J.L.R., Griffanti, L., Douaud, G., Okell, T.W., Weale, P., Dragou, I., Garratt, S., Hudson, S., Collins, R., Jenkinson, M., Matthews, P.M., Smith, S.M., 2016. Multimodal population brain imaging in the UK biobank prospective epidemiological study. *Nat. Neurosci.* 19, 1523–1536. doi:10.1038/nn.4393.
- Mitchell, T.J., Beauchamp, J.J., 1988. Bayesian variable selection in linear regression. *J. Am. Stat. Assoc.* 83 (404), 1023–1032. doi:10.1080/01621459.1988.10478694.
- Mueller, S., Wang, D., Fox, M., Yeo, B.T., Sepulcre, J., Sabuncu, M., Shafee, R., Lu, J., Liu, H., 2013. Individual variability in functional connectivity architecture of the human brain. *Neuron* 77 (3), 586–595. doi:10.1016/j.neuron.2012.12.028.
- Nickerson, L.D., Smith, S.M., Öngür, D., Beckmann, C.F., 2017. Using dual regression to investigate network shape and amplitude in functional connectivity analyses. *Front Neurosci.* 11 (115). doi:10.3389/fnins.2017.00115.
- Noble, K.G., Houston, S.M., Brito, N.H., Bartsch, H., Kan, E., Kuperman, J.M., Akshoomoff, N., Amaral, D.G., Bloss, C.S., Libiger, O., Schork, J.L.R., Murray, S.S., Casey, B.J., Chang, L., Ernst, T.M., Frazier, J.A., Gruen, J.R., Kennedy, D.N., Van Zijl, P., Mostofsky, S., Kaufmann, W.E., Kenet, T., Dale, A.M., Jernigan, T.L., Sowell, E.R., 2015. Family income, parental education and brain structure in children and adolescents. *Nat. Neurosci.* 18, 773–778. doi:10.1038/nn.3983.
- Noble, S., Scheinost, D., Finn, E.S., Shen, X., Papademetris, X., McEwen, S.C., Bearden, C.E., Addington, J., Goodyear, B., Cadenhead, K.S., Mirzakhani, H., Cornblatt, B.A., Olvet, D.M., Mathalon, D.H., McGlashan, T.H., Perkins, D.O., Belger, A., Seidman, L.J., Thermenos, H., Tsuang, M.T., van Erp, T.G.M., Walker, E.F., Hamann, S., Woods, S.W., Cannon, T.D., Constable, R.T., 2017. Multisite reliability of MR-based functional connectivity. *NeuroImage* 146, 959–970. doi:10.1016/j.neuroimage.2016.10.020.
- Pervais, U., Vidaurre, D., Woolrich, M.W., Smith, S.M., 2020. Optimising network modelling methods for fmri. *NeuroImage* 211, 116604. doi:10.1016/j.neuroimage.2020.116604.
- Poldrack, R.A., Barch, D., Mitchell, J., Wager, T., Wagner, A., Devlin, J., Cumba, C., Koyejo, O., Milham, M., 2013. Toward open sharing of task-based fmri data: the openfmri project. *Front Neuroinform* 7 (12). doi:10.3389/fninf.2013.00012.



- Poldrack, R.A., Laumann, T.O., Koyejo, O., Gregory, B., Hover, A., Chen, M.-Y., Gorgolewski, K.J., Luci, J., Joo, S.J., Boyd, R.L., Hunnicke-Smith, S., Simpson, Z.B., Caven, T., Sochat, V., Shine, J.M., Gordon, E., Snyder, A.Z., Adeyemo, B., Petersen, S.E., Glahn, D.C., Reese McKay, D., Curran, J.E., Gring, H.H.H., Carless, M.A., Blangero, J., Dougherty, R., Leemans, A., Handwerker, D.A., Frick, L., Marcotte, E.M., Mumford, J.A., 2015. Long-term neural and physiological phenotyping of a single human. *Nat. Commun.* 6 (8885). doi:10.1038/ncomms9885.
- Power, J.D., Barnes, K.A., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., 2012. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *NeuroImage* 59 (3), 2142–2154. doi:10.1016/j.neuroimage.2011.10.018.
- Power, J.D., Plitt, M., Laumann, T.O., Martin, A., 2017. Sources and implications of whole-brain fMRIsignals in humans. *NeuroImage* 146, 609–625. doi:10.1016/j.neuroimage.2016.09.038.
- Preti, M.G., Van De Ville, D., 2017. Dynamics of functional connectivity at high spatial resolution reveal long-range interactions and fine-scale organization. *Sci. Rep.* 7 (1), 12773. doi:10.1038/s41598-017-12993-1.
- Qing, Z., Gong, G., 2016. Size matters to function: brain volume correlates with intrinsic brain activity across healthy individuals. *NeuroImage* 139, 271–278. doi:10.1016/j.neuroimage.2016.06.046.
- Raemaekers, M., Schellekens, W., van Wezel, R.J.A., Petridou, N., Kristo, G., Ramsey, N.F., 2014. Patterns of resting state connectivity in human primary visual cortical areas: a 7 t fmri study. *NeuroImage* 84, 911–921. doi:10.1016/j.neuroimage.2013.09.060.
- Raichle, M.E., MacLeod, A.M., Snyder, A.Z., Powers, W.J., Gusnard, D.A., Shulman, G.L., 2001. A default mode of brain function. *Proceedings of the National Academy of Sciences* 98 (2), 676–682. doi:10.1073/pnas.98.2.676.
- Reiss, A.L., Abrams, M.T., Singer, H.S., Ross, J.L., Denckla, M.B., 1996. Brain development, gender and IQ in children: avolumetric imaging study. *Brain* 119 (5), 1763–1774. doi:10.1093/brain/119.5.1763.
- Robert, P., Escoufier, Y., 1976. A unifying tool for linear multivariate statistical methods: the RV-coefficient. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 25 (3), 257–265.
- Robinson, E.C., Garcia, K., Glasser, M.F., Chen, Z., Coalson, T.S., Makropoulos, A., Bozcek, J., Wright, R., Schuh, A., Webster, M., Hutter, J., Price, A., Cordero Grande, L., Hughes, E., Tumor, N., Bayly, P.V., Van Essen, D.C., Smith, S.M., Edwards, A.D., Hajnal, J., Jenkinson, M., Glocker, B., Rueckert, D., 2018. Multimodal surface matching with higher-order smoothness constraints. *NeuroImage* 167, 453–465. doi:10.1016/j.neuroimage.2017.10.037.
- Robinson, E.C., Jbabdi, S., Glasser, M.F., Andersson, J., Burgess, G.C., Harms, M.P., Smith, S.M., Van Essen, D.C., Jenkinson, M., 2014. MSM: A new flexible framework for multimodal surface matching. *NeuroImage* 100, 414–426. doi:10.1016/j.neuroimage.2014.05.069.
- Sala-Llonch, R., Smith, S.M., Woolrich, M.W., Duff, E.P., 2019. Spatial parcellations, spectral filtering, and connectivity measures in fMRI: optimizing for discrimination. *Hum. Brain Mapp.* 40 (2), 407–419. doi:10.1002/hbm.24381.
- Salehi, M., Greene, A.S., Karbasi, A., Shen, X., Scheinost, D., Constable, R.T., 2020. There is no single functional atlas even for a single individual: functional parcel definitions change with task. *NeuroImage* 208, 116366. doi:10.1016/j.neuroimage.2019.116366.
- Salehi, M., Karbasi, A., Scheinost, D., Constable, R.T., 2017. A submodular approach to create individualized parcellations of the human brain. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (Eds.), *Medical Image Computing and Computer Assisted Intervention—MICCAI 2017*. Springer International Publishing, Cham, pp. 478–485.
- Satterthwaite, T.D., Wolf, D.H., Loughead, J., Ruparel, K., Elliott, M.A., Hakonarson, H., Gur, R.C., Gur, R.E., 2012. Impact of in-scanner head motion on multiple measures of functional connectivity: relevance for studies of neurodevelopment in youth. *NeuroImage* 60 (1), 623–632. doi:10.1016/j.neuroimage.2011.12.063.
- Scott, A., Courtney, W., Wood, D., De la Garza, R., Lane, S., Wang, R., King, M., Roberts, J., Turner, J., Calhoun, V.D., 2011. COINS: An innovative informatics and neuroimaging tool suite built for large heterogeneous datasets. *Front. Neuroinform.* 5 (33). doi:10.3389/fninf.2011.00033.
- Shackman, A.J., Salomons, T.V., Slagter, H.A., Fox, A.S., Winter, J.J., Davidson, R.J., 2011. The integration of negative affect, pain and cognitive control in the cingulate cortex. *Nat. Rev. Neurosci.* 12 (3), 154–167. doi:10.1038/nrn2994.
- Shaw, P., Greenstein, D., Lerch, J., Clasen, L., Lenroot, R., Gogtay, N., Evans, A., Rapoport, J., Giedd, J., 2006. Intellectual ability and cortical development in children and adolescents. *Nature* 440, 676–679. doi:10.1038/nature04513.
- Shehzad, Z., Kelly, A.M.C., Reiss, P.T., Gee, D.G., Gotimer, K., Uddin, L.Q., Lee, S.H., Margulies, D.S., Roy, A.K., Biswal, B.B., Petkova, E., Castellanos, F.X., Milham, M.P., 2009. The resting brain: unconstrained yet reliable. *Cerebral Cortex* 19 (10), 2209–2229. doi:10.1093/cercor/bhn256.
- Shirer, W.R., Ryali, S., Rykhlevskaia, E., Menon, V., Greicius, M.D., 2011. Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cerebral Cortex* 22 (1), 158–165. doi:10.1093/cercor/bhr099.
- Shmuel, A., Yacoub, E., Chaimov, D., Logothetis, N.K., Ugurbil, K., 2007. Spatio-temporal point-spread function of fMRI signal in human gray matter at 7 tesla. *NeuroImage* 35 (2), 539–552. doi:10.1016/j.neuroimage.2006.12.030.
- Shulman, G.L., Fiez, J.A., Corbetta, M., Buckner, R.L., Miezin, F.M., Raichle, M.E., Petersen, S.E., 1997. Common blood flow changes across visual tasks: I. decreases in cerebral cortex. *J. Cogn. Neurosci.* 9 (5), 648–663. doi:10.1162/jocn.1997.9.5.648.
- Smith, S.M., Feingberg, D.A., Griffanti, L., Harms, M.P., Kelly, M., Laumann, T.O., Miller, K.L., Moeller, S., Petersen, S., Power, J., Salimi-Khorshidi, G., Snyder, A.Z., Vu, A.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. doi:10.1016/j.neuroimage.2013.05.039. Mapping the Connectome.
- Smith, S.M., Hyvärinen, A., Varoquaux, G., Miller, K.L., Beckmann, C.F., 2014. Group-PCA for very large fmri datasets. *NeuroImage* 101, 738–749. doi:10.1016/j.neuroimage.2014.07.051.
- 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., Feingberg, D.A., Yacoub, E.S., Ugurbil, K., 2012. Temporally-independent functional modes of spontaneous brain activity. *Proceedings of the National Academy of Sciences* 109 (8), 3131–3136. doi:10.1073/pnas.1121329109.
- Smith, S.M., Miller, K.L., Salimi-Khorshidi, G., Webster, M., Beckmann, C.F., Nichols, T.E., Ramsey, J.D., Woolrich, M.W., 2011. Network modelling methods for FMRI. *NeuroImage* 54 (2), 875–891. doi:10.1016/j.neuroimage.2010.08.063.
- Smith, S.M., Nichols, T.E., Vidaurre, D., Winkler, A.M., Behrens, T.E.J., Glasser, M.F., Ugurbil, K., Barch, D.M., Van Essen, D.C., Miller, K.L., 2015. A positive-negative mode of population covariation links brain connectivity, demographics and behavior. *Nat. Neurosci.* 18 (11), 1565–1567. doi:10.1038/nn.4125.
- Smith, S.M., Vidaurre, D., Beckmann, C.F., Glasser, M.F., Jenkinson, M., Miller, K.L., Nichols, T.E., Robinson, E.C., Salimi-Khorshidi, G., Woolrich, M.W., Barch, D.M., Ugurbil, K., Van Essen, D.C., 2013. Functional connectomics from resting-state fmri. *Trends Cogn. Sci. (Regul. Ed.)* 17 (12), 666–682.
- Stegle, O., Lippert, C., Mooij, J., Lawrence, N., Borgwardt, K., 2011. Efficient inference in matrix-variate Gaussian models with iid observation noise. In: *Shawe-Taylor, J., Zemel, R., Bartlett, P., Pereira, F., Weinberger, K. (Eds.), Advances in Neural Information Processing Systems 24*, pp. 630–638. Scientific Publication Twenty-Fifth Annual Conference on Neural Information Processing Systems, 12–17 December, 2011 Granada (Spain).
- Stein, J.L., Medland, S.E., Vasquez, A.A., Hibar, D.P., Senstad, R.E., Winkler, A.M., Toro, R., Appel, K., Bartcecek, R., Bergmann, O., Bernard, M., Brown, A.A., Cannon, D.M., Chakravarty, M.M., Christoforou, A., Domin, M., Grimm, O., Hollinshead, M., Holmes, A.J., Homuth, G., Hottenga, J.-J., Langan, C., Lopez, L.M., Hansell, N.K., Hwang, K.S., Kim, S., Laje, G., Lee, P.H., Liu, X., Loth, E., Lourdasamy, A., Mattingsdal, M., Mohnke, S., Maniega, S.M., Nho, K., Nugent, A.C., O'Brien, C., Pappmeyer, M., Putz, B., Ramasamy, A., Rasmussen, J., Rijpkema, M., Risacher, S.L., Roddey, J.C., Rose, E.J., Ryten, M., Shen, L., Sprooten, E., Strengman, E., Teumer, A., Trabzuni, D., Turner, J., van Eijk, K., van Erp, T.G.M., van Tol, M.-J., Wittfeld, K., Wolf, C., Woudstra, S., Aleman, A., Alhusaini, S., Almsay, L., Binder, E.B., Brohawn, D.G., Cantor, R.M., Carless, M.A., Corvin, A., Czisch, M., Curran, J.E., Davies, G., de Almeida, M.A.A., Delanty, N., Depondt, C., Duggirala, R., Dyer, T.D., Erk, S., Fagerness, J., Fox, P.T., Freimer, N.B., Gill, M., Goring, H.H.H., Hagler, D.J., Hoehn, D., Holsboer, F., Hoogman, M., Hosten, N., Jahanshad, N., Johnson, M.P., Kasperaviciute, D., Kent, J.W., Kochunov, P., Lancaster, J.L., Lawrie, S.M., Liewald, D.C., Mandl, R., Matarin, M., Mattheisen, M., Meisenzahl, E., Melle, I., Moses, E.K., Muhleisen, T.W., Nauck, M., Nothen, M.M., Olvera, R.L., Pandolfo, M., Pike, G.B., Puls, R., Reinvang, I., Renteria, M.E., Rietschel, M., Roffman, J.L., Royle, N.A., Rujescu, D., Savitz, J., Schnack, H.G., Schnell, K., Seiferth, N., Smith, C., Steen, V.M., Valdes Hernandez, M.C., Van den Heuvel, M., van der Wee, N.J., Van Haren, N.E.M., Veltman, J.A., Volzke, H., Walker, R., Westlye, L.T., Whelan, C.D., Agartz, I., Boomsma, D.I., Cavalleri, G.L., Dale, A.M., Djurovic, S., Drevets, W.C., Hagoot, P., Hall, J., Heinz, A., Jack, C.R., Foroud, T.J., Le Hellard, S., Macciardi, F., Montgomery, G.W., Poline, J.B., Porteous, D.J., Sisodiya, S.M., Starr, J.M., Sussmann, J., Toga, A.W., Veltman, D.J., Walter, H., Weiner, M.W., the Alzheimer's Disease Neuroimaging Initiative (ADNI), EPICgen consortium, IMAGEN consortium, Saguayan Youth Study Group (SYS), Bis, J.C., Ikram, M.A., Smith, A.V., Gudnason, V., Tzourio, C., Vernooij, M.W., Launer, L.J., DeCarli, C., Seshadri, S., Cohorts for Heart and Aging Research in Genomic Epidemiology (CHARGE) Consortium, Andreassen, O.A., Apostolova, L.G., Bastin, M.E., Blangero, J., Brunner, H.G., Buckner, R.L., Cichon, S., Coppola, G., de Zubicaray, G.I., Deary, I.J., Donohoe, G., de Geus, E.J.C., Espeseth, T., Fernández, G., Glahn, D.C., Grabe, H.J., Hardy, J., Hulshoff Pol, H.E., Jenkinson, M., Kahn, R.S., McDonald, C., McIntosh, A.M., McMahon, F.J., McEhonor, K.L., Meyer-Lindenberg, A., Morris, D.W., Müller-Yamshok, B., Nichols, T.E., Ophoff, R.A., Paus, T., Pausova, Z., Penninx, B.W., Potkin, S.G., Sämann, P.G., Saykin, A.J., Schumann, G., Smoller, J.W., Wardlaw, J.M., Weale, M.E., Martin, N.G., Franke, B., Wright, M.J., Thompson, P.M., 2012. Identification of common variants associated with human hippocampal and intracranial volumes. *Nat. Genet.* 44 (5), 552–561. doi:10.1038/ng.2250.
- Stephan, K., Iglesias, S., Heinzle, J., Diaconescu, A., 2015. Translational perspectives for computational neuroimaging. *Neuron* 87 (4), 716–732. doi:10.1016/j.neuron.2015.07.008.
- Stephan, K.E., Schlaghauf, F., Huys, Q.J.M., Raman, S., Aponte, E.A., Brodersen, K.H., Rigoux, L., Moran, R.J., Daunizeau, J., Dolan, R.J., Friston, K.J., Heinz, A., 2017. Computational neuroimaging strategies for single patient predictions. *NeuroImage* 145, Part B, 180–199. doi:10.1016/j.neuroimage.2016.06.038. Individual Subject Prediction.
- Tagliazucchi, E., Laufs, H., 2014. Decoding wakefulness levels from typical fmri resting-state data reveals reliable drifts between wakefulness and sleep. *Neuron* 82 (3), 695–708. doi:10.1016/j.neuron.2014.03.020.
- Thompson, P.M., Stein, J.L., Medland, S.E., Hibar, D.P., Vasquez, A.A., Renteria, M.E., Toro, R., Jahanshad, N., Schumann, G., Franke, B., Wright, M.J., Martin, N.G., Agartz, I., Alda, M., Alhusaini, S., Almsay, L., Almeida, J., Alpert, K., Andreassen, N.C., Andreassen, O.A., Apostolova, L.G., Appel, K., Armstrong, N.J., Arribas, B., Bastin, M.E., Bauer, M., Bearden, C.E., Bergmann, Ø., Binder, E.B., Blangero, J., Bockholt, H.J., Bøen, E., Bois, C., Boomsma, D.I., Booth, T., Bowman, I.J., Bralten, J., Brouwer, R.M., Brunner, H.G., Brohawn, D.G., Buckner, R.L., Buitelaar, J., Bulayeva, K., Bustillo, J.R., Calhoun, V.D., Cannon, D.M., Cantor, R.M., Carless, M.A., Caseras, X., Cavalleri, G.L., Chakravarty, M.M., Chang, K.D., Ching, C.R.K., Christoforou, A., Cichon, S., Clark, V.P., Conrod, P., Coppola, G., Crespo-Facorro, B., Curran, J.E., Czisch, M., Deary, I.J., de Geus, E.J.C., de Graaf, A., Delvecchio, G., Depondt, C., de Haan, L., de Zubicaray, G.I., Dima, D., Dimitrova, R., Djurovic, S.,

- Dong, H., Donohoe, G., Duggirala, R., Dyer, T.D., Ehrlich, S., Ekman, C.J., Elvåshagen, T., Emsell, L., Erk, S., Espeseth, T., Fagerness, J., Fears, S., Fedko, I., Fernández, G., Fisher, S.E., Foroud, T., Fox, P.T., Francks, C., Frangou, S., Frey, E.M., Frodl, T., Frouin, V., Garavan, H., Giddaluru, S., Glahn, D.C., Godlewska, B., Goldstein, R.Z., Gollub, R.L., Graber, H.J., Grimm, O., Gruber, O., Guadalupe, T., Gur, R.E., Gur, R.C., Göring, H.H.H., Hagenaars, S., Hajek, T., Hall, G.B., Hall, J., Hardy, J., Hartman, C.A., Hass, J., Hatton, S.N., Haukvik, U.K., Hegenscheid, K., Heinz, A., Hickie, I.B., Ho, B.-C., Hoehn, D., Hoekstra, P.J., Hollinshead, M., Holmes, A.J., Homuth, G., Hoogman, M., Hong, L.E., Hosten, N., Hottenga, J.-J., Hulshoff Pol, H.E., Hwang, K.S., Jack Jr., C.R., Jenkinson, M., Johnston, C., Jönsson, E.G., Kahn, R.S., Kasperaviciute, D., Kelly, S., Kim, S., Kochunov, P., Koenders, L., Krämer, B., Kwok, J.B.J., Lagopoulos, J., Laje, G., Landen, M., Landman, B.A., Lauriello, J., Lawrie, S.M., Lee, P.H., Le Hellard, S., Lemaître, H., Leonardo, C.D., Li, C.-s., Liberg, B., Liewald, D.C., Liu, X., Lopez, H.M., Luth, E., Lourdasamy, A., Luciano, M., Macciardi, F., Machielsen, M.W.J., MacQueen, G.M., Malt, U.F., Mandl, R., Manoach, D.S., Martinot, J.-L., Matarin, M., Mather, K.A., Mattheisen, M., Mattingsdal, M., Meyer-Lindenberg, A., McDonald, C., McIntosh, A.M., McMahon, F.J., McMahon, K.L., Meisenzahl, E., Melle, I., Milaneschi, Y., Mohnke, S., Montgomery, G.W., Morris, D.W., Moses, E.K., Mueller, B.A., Muñoz Maniega, S., Mühleisen, T.W., Müller-Myhsok, B., Mwangi, B., Nauck, M., Nho, K., Nichols, T.E., Nilsson, L.-G., Nugent, A.C., Nyberg, L., Olvera, R.L., Oosterlaan, J., Ophoff, R.A., Pandolfo, M., Papalampropoulou-Tsiridou, M., Pappmeyer, M., Paus, T., Pausova, Z., Pearson, G.D., Penninx, B.W., Peterson, C.P., Pfennig, A., Phillips, M., Pike, G.B., Poline, J.-B., Potkin, S.G., Pütz, B., Ramasamy, A., Rasmussen, J., Rietschel, M., Rijpkema, M., Risacher, S.L., Roffman, J.L., Roiz-Santiañez, R., Romanczuk-Seiferth, N., Rose, E.J., Royle, N.A., Rujescu, D., Ryten, M., Sachdev, P.S., Salami, A., Satterthwaite, T.D., Savitz, J., Saykin, A.J., Scanlon, C., Schmaal, L., Schnack, H.G., Schork, A.J., Schulz, S.C., Schür, R., Seidman, L., Shen, L., Shoemaker, J.M., Simmons, A., Sisodiya, S.M., Smith, C., Smoller, J.W., Soares, J.C., Sponheim, S.R., Sprooten, E., Starr, J.M., Steen, V.M., Strakowski, S., Strike, L., Sussmann, J., Sämann, P.G., Teumer, A., Toga, A.W., Tordesillas-Gutierrez, D., Trabzuni, D., Trost, S., Turner, J., Van den Heuvel, M., van der Wee, N.J., van Eijk, K., van Erp, T.G.M., van Haren, N.E.M., van 't Ent, D., van Tol, M.-J., Valdés Hernández, M.C., Veltman, D.J., Versace, A., Völzke, H., Walker, R., Walter, H., Wang, L., Wardlaw, J.M., Weale, M.E., Weiner, M.W., Wen, W., Westlye, L.T., Whalley, H.C., Whelan, C.D., White, T., Winkler, A.M., Wittfeld, K., Woldehawariat, G., Wolf, C., Zilles, D., Zwiwers, M.P., Thalamuthu, A., Schofield, P.R., Freimer, N.B., Lawrence, N.S., Drevets, W., 2014. The enigma consortium: large-scale collaborative analyses of neuroimaging and genetic data. *Brain Imaging Behav.* 8 (2), 153–182. doi:10.1007/s11682-013-9269-5.
- Titsias, M., Lázaro-Gredilla, M., 2011. Spike and Slab Variational Inference for Multi-task and Multiple Kernel Learning. In: Shawe-Taylor, J., Zemel, R., Bartlett, P., Pereira, F., Weinberger, K. (Eds.), *Advances in Neural Information Processing Systems 24*, pp. 2339–2347.
- Tong, T., Aganj, I., Ge, T., Polimeni, J.R., Fischl, B., 2017. Functional density and edge maps: characterizing functional architecture in individuals and improving cross-subject registration. *Neuroimage* 158, 346–355. doi:10.1016/j.neuroimage.2017.07.019.
- Van Dijk, K.R.A., Sabuncu, M.R., Buckner, R.L., 2012. The influence of head motion on intrinsic functional connectivity mri. *NeuroImage* 59 (1), 431–438. doi:10.1016/j.neuroimage.2011.07.044. Neuroergonomics: The human brain in action and at work.
- Van Essen, D.C., Dierker, D.L., 2007. Surface-based and probabilistic atlases of primate cerebral cortex. *Neuron* 56 (2), 209–225. doi:10.1016/j.neuron.2007.10.015.
- Van Essen, D.C., Glasser, M.F., Dierker, D.L., Harwell, J., Coalson, T., 2012. Parcellations and hemispheric asymmetries of human cerebral cortex analyzed on surface-based atlases. *Cerebral Cortex* 22 (10), 2241–2262. doi:10.1093/cercor/bhr291.
- Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E.J., Yacoub, E., Ugurbil, K., 2013. The WU-Minn Human Connectome Project: An overview. *NeuroImage* 80, 62–79. doi:10.1016/j.neuroimage.2013.05.041. Mapping the Connectome.
- Van Essen, D.C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T.E.J., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S., Penna, S.D., Feinberg, D., Glasser, M.F., Harel, N., Heath, A., Larson-Prior, L., Marcus, D., Michalareas, G., Moeller, S., Oostenveld, R., Petersen, S., Prior, F., Schlaggar, B., Smith, S.M., Snyder, A., Xu, J., Yacoub, E., 2012. The Human Connectome Project: A data acquisition perspective. *NeuroImage* 62 (4), 2222–2231. doi:10.1016/j.neuroimage.2012.02.018. Connectivity.
- Vanderwal, T., Eilbott, J., Finn, E.S., Craddock, R.C., Turnbull, A., Castellanos, F.X., 2017. Individual differences in functional connectivity during naturalistic viewing conditions. *Neuroimage* 157, 521–530. doi:10.1016/j.neuroimage.2017.06.027.
- Varoquaux, G., Gramfort, A., Pedregosa, F., Michel, V., Thirion, B., 2011. Multi-subject Dictionary Learning to Segment an Atlas of Brain Spontaneous Activity. In: Székely, G., Hahn, H.K. (Eds.), *Information Processing in Medical Imaging*. In: *Lecture Notes in Computer Science*, 6801. Springer Berlin Heidelberg, pp. 562–573. doi:10.1007/978-3-642-22092-0\_46.
- Varoquaux, G., Raamana, P.R., Engemann, D.A., Hoyos-Idrobo, A., Schwartz, Y., Thirion, B., 2017. Assessing and tuning brain decoders: Cross-validation, caveats, and guidelines. *NeuroImage* 145, Part B, 166–179. doi:10.1016/j.neuroimage.2016.10.038. Individual Subject Prediction.
- Varoquaux, G., Sadaghiani, S., Pinel, P., Kleinschmidt, A., Poline, J., Thirion, B., 2010. A group model for stable multi-subject ICA on fmri datasets. *Neuroimage* 51 (1), 288–299. doi:10.1016/j.neuroimage.2010.02.010.
- Vidaurre, D., Smith, S.M., Woolrich, M.W., 2017. Brain network dynamics are hierarchically organized in time. *Proceedings of the National Academy of Sciences* 114 (48), 12827–12832. doi:10.1073/pnas.1705120114.
- Wang, D., Buckner, R.L., Fox, M.D., Holt, D.J., Holmes, A.J., Stoecklein, S., Langs, G., Pan, R., Qian, T., Li, K., Baker, J.T., Stufflebeam, S.M., Wang, K., Wang, X., Hong, B., Liu, H., 2015. Parcellating cortical functional networks in individuals. *Nat. Neurosci.* 18 (12), 1853–1860. doi:10.1038/nn.4164.
- Winkler, A.M., Ridgway, G.R., Douaud, G., Nichols, T.E., Smith, S.M., 2016. Faster permutation inference in brain imaging. *Neuroimage* 141, 502–516. doi:10.1016/j.neuroimage.2016.05.068.
- Winkler, A.M., Ridgway, G.R., Webster, M.A., Smith, S.M., Nichols, T.E., 2014. Permutation inference for the general linear model. *Neuroimage* 92, 381–397. doi:10.1016/j.neuroimage.2014.01.060.
- Winn, J., Bishop, C.M., 2005. Variational message passing. *Journal of Machine Learning Research* 6, 661–694.
- Woolrich, M.W., Behrens, T.E.J., Smith, S.M., 2004. Constrained linear basis sets for HRF modelling using variational bayes. *Neuroimage* 21 (4), 1748–1761. doi:10.1016/j.neuroimage.2003.12.024.
- Worsley, K.J., Evans, A.C., Marrett, S., Neelin, P., 1996. A three-dimensional statistical analysis for CBF activation studies in human brain. *Journal of Cerebral Blood Flow & Metabolism* 12 (6), 900–918. doi:10.1038/jcbfm.1992.127.
- Zang, Y.-F., He, Y., Zhu, C.-Z., Cao, Q.-J., Sui, M.-Q., Liang, M., Tian, L.-X., Jiang, T.-Z., Wang, Y.-F., 2007. Altered baseline brain activity in children with ADHD revealed by resting-state functional MRI. *Brain and Development* 29 (2), 83–91. doi:10.1016/j.braindev.2006.07.002.
- Zöllner, D.M., Bolton, T.A.W., Karahanoglu, F.I., Eliez, S., Schaefer, M., Van De Ville, D., 2019. Robust recovery of temporal overlap between network activity using transient-informed spatio-temporal regression. *IEEE Trans. Med. Imaging* 38 (1), 291–302. doi:10.1109/TMI.2018.2863944.
- Zou, Q.-H., Zhu, C.-Z., Yang, Y., Zuo, X.-N., Long, X.-Y., Cao, Q.-J., Wang, Y.-F., Zang, Y.-F., 2008. An improved approach to detection of amplitude of low-frequency fluctuation (ALFF) for resting-state fmri: fractional ALFF. *J. Neurosci. Methods* 172 (1), 137–141. doi:10.1016/j.jneumeth.2008.04.012.
- Zuo, X.-N., Di Martino, A., Kelly, C., Shehzad, Z.E., Gee, D.G., Klein, D.F., Castellanos, F.X., Biswal, B.B., Milham, M.P., 2010. The oscillating brain: complex and reliable. *Neuroimage* 49 (2), 1432–1445. doi:10.1016/j.neuroimage.2009.09.037.
- Zuo, X.-N., Kelly, C., Adelstein, J.S., Klein, D.F., Castellanos, F.X., Milham, M.P., 2010. Reliable intrinsic connectivity networks: test-retest evaluation using ICA and dual regression approach. *Neuroimage* 49 (3), 2163–2177. doi:10.1016/j.neuroimage.2009.10.080.