Directed functional connectivity using dynamic graphical models

Simon Schwab[1,2,3,*], Ruth Harbord[4], Valerio Zerbi[5], Lloyd Elliott[6], Soroosh Afyouni[2], Jim Q. Smith[3], Mark W. Woolrich[7], Stephen M. Smith[7], Thomas E. Nichols[1,2,3]

1 Big Data Institute, Li Ka Shing Centre for Health Information and Discovery, Nuffield Department of Population Health, University of Oxford, United Kingdom

2 Institute of Digital Healthcare, WMG, University of Warwick, United Kingdom

3 Department of Statistics, University of Warwick, United Kingdom

4 MOAC Doctoral Training Centre, University of Warwick, United Kingdom

5 Neural Control of Movement Lab, Department of Health Sciences and Technology, ETH Zurich, Switzerland

6 Department of Statistics, University of Oxford, United Kingdom

7 Oxford Centre for Functional MRI of the Brain, University of Oxford, Oxford, United Kingdom

* Corresponding author: Simon Schwab, Big Data Institute, Li Ka Shing Centre for Health Information and Discovery, University of Oxford, Old Road Campus, Oxford OX3 7LF, United Kingdom, simon.schwab@bdi.ac.ox.uk

Funding: Simon Schwab acknowledges funding from the Swiss National Science Foundation (SNSF, No. 162066). VZ is supported by ETH Career Seed Grant SEED-42 16-1. RH acknowledges support from an EPSRC PhD studentship (grant number EP/F500378/1). TN is supported by the Wellcome Trust, 100309/Z/12/Z.

Data Availability Statement: Data with reproducible analyses can be found at https://github.com/schw4b/DGM. The source code of the R package “multdyn” is available at https://github.com/schw4b/multdyn.

Abstract

There are a growing number of neuroimaging methods that model spatio-temporal patterns of brain activity to allow more meaningful characterizations of brain networks. However, directed relationships in networks are difficult to estimate, and only very few methods are available. Here, we consider a method for dynamic, directed functional connectivity: Dynamic graphical models (DGM) are a multivariate graphical model with dynamic, time-varying coefficients that describe instantaneous directed relationships between nodes. Parameter estimation and model selection are implemented by using Bayesian methods. A further benefit of DGM is that networks may contain loops. We use network simulations, human resting-state-fMRI (N = 500) and mouse fMRI (N = 16) to investigate the validity and reliability of the estimated networks. We simulated systematic lags of the hemodynamic response at different brain regions to investigate how these potentially bias directionality estimates. In the presence of such lag confounds (0.4–0.8 seconds offset between connected nodes), our method demonstrated sensitivity of 72%–77% and a specificity of 67% to detect the true direction. Stronger lag confounds reduced sensitivity, but did not increase false positives (i.e., directionality estimates of the opposite direction). On real data, we investigated the reliability of network estimates across and within subjects and found the DMN having inputs to cerebellar and limbic networks, as well as the reciprocal relationship between the visual medial and visual lateral network to be the most consistent relationships. Finally, we tested the method in mouse fMRI data and discovered a highly plausible relationship between areas in the hippocampus feeding to the cingulate cortex. We provide a computationally efficient implementation of DGM as a free software package (“multdyn”) for R.

Key words: directed dynamic connectivity; effective connectivity; dynamic graphical models; time-varying connectivity; resting-state fMRI;

Introduction

Human behavior is underpinned by brain circuits and there is evidence of abnormal connectivity in brain networks related to the psychopathology of mental disorders (Bassett and Bullmore, 2009; Buckner et al., 2008; Menon, 2011), for example, in schizophrenia (Zhou et al., 2016), major depression (Cheng et al., 2016; Guo et al., 2016), dementia (Chase, 2014), or anxiety disorders (Peterson et al., 2014). New treatment strategies in the the future may benefit from new computational method developments that lead to an improved understanding of altered functional networks of the brain (Friston et al., 2014b; Huys et al., 2016; Stephan and Mathys, 2014). However, a majority of studies in functional connectivity (FC) are based on correlations (full, or partial) between fMRI time series, and even though they are fast to compute and can handle large networks, these have some major limitations (Smith et al., 2013b).

First, FC in terms of correlation between time series reflects static, stationary relationships with fixed connection strengths across time. However, coupling between brain systems is not constant over time—there is evidence that brain connectivity is better described in terms of time-varying connectivity (Allen et al., 2014; Liu and Duyn, 2013; Shine and Poldrack, 2017; Smith et al., 2012b). Emerging evidence suggests that dynamic FC may indicate changes in macroscopic activity patterns underlying critical aspects of cognition and behavior (Allen et al., 2014; Braun et al., 2015; Hutchison et al., 2013a; Smith et al., 2012a; Zalesky et al., 2014), which can be highly relevant in psychiatry as well (Damaraju et al., 2014; Kaiser et al., 2016; Zhang et al., 2016). Such dynamic coupling can occur in different behavioral contexts and during cognitive control (Cocchi et al., 2013; Gonzalez-Castillo et al., 2015; Shine et al., 2016). In particular, intermodular connections seem to be dynamic, linking different systems of the brain (Zalesky et al., 2014). Hence, it is more plausible that the brain exhibits dynamic communication (Andrews-Hanna et al., 2014; Betzel et al., 2016; Bressler and Kelso, 2001; Fox et al., 2005; Hutchison et al., 2013b). Only recently have studies begun to look closer at fluctuating dynamic connectivity, but proper validation is important (Laumann et al., 2016; Leonardi and Van De Ville, 2015; Lindquist et al., 2014; Ryali et al., 2011). Consequently, dynamic methods will potentially have greater impact than static methods in the future (Bassett and Gazzaniga, 2011; Calhoun et al., 2014; Preti et al., 2016; Vidaurre et al., 2017).

Second, FC in terms of a correlation between nodes reflects undirected relationships between them. A major challenge is to estimate the direction of information flow in networks (Ramsey et al., 2010). Notably, the hemodynamic response (HR) is subjected some temporal variation. Even though most variation is between subjects, there is also variation between brain regions (Handwerker et al., 2004) which can be problematic. As a result, lag-based methods, such as Granger causality, can be confounded by the variability of the HR in different brain areas (Smith et al., 2013b, 2011). For example, if a node A is influencing a node B, but node A has a slower HR compared to B, lag-based methods may estimate the wrong direction. Another popular method to estimate directed networks is dynamic causal modeling (DCM; Friston et al., 2003). However, DCM cannot efficiently search across all possible network models, and thus prior knowledge is needed to constrain the model space. Only small networks are being investigated in DCM studies, often less than 10 nodes (Hein et al., 2016; Nakataki et al., 2017). Also, the standard form of DCM assumes fixed connection strengths and, as initially designed for task fMRI, is not suitable for resting-state fMRI (Stephan et al., 2010). Newer work on stochastic DCM (Friston et al., 2014a; Razi et al., 2015) and regression DCM (rDCM; Frässle et al., 2017) have sought to address these limitations. rDCMs can estimate larger networks but the effect of the HR variability in different regions is a limitation acknowledged in that work.

In the present study, we use dynamic graphical models (DGM) to estimate dynamic, directed functional connectivity. It is plausible that brain areas are not acyclic due to its reciprocal polysynaptic connections (Friston, 2011). Therefore, unlike our related work on “multiregression dynamic models” (MDMs; Costa et al., 2015), we do not constrain the networks to be acyclic. Without the acyclic constraint, we refer to our approach as the Dynamic Graphical Model (DGM). DGM are a dynamic network with time-varying connection strengths that can accommodate larger directed networks than DCM. Further, DGM have added interpretability as a regression model: for each child node, its time series are regressed on the time series from one or more other parent nodes; the directed relationship corresponds to the information flowing from the parent to the child node. Critically, the DGM use instantaneous interactions between the brain regions rather than lagged relationships to estimate directionality.

A major challenge of DGM is the number of possible graphs which increase exponentially as a function of the number of nodes. Searching such a large solution space is not feasible, even with modern

computational resources. For example, for R = 15 regions/nodes there $2^{n^2-n} \approx 10^{63}$ possible directed graphs[1] . However, the DGM model evidence factorizes by node (see below), and hence models can be optimised for each node independently, and thus only $n2^{n-1}$ model evaluations (e.g. 245,760 for R = 15) are required. Even so, networks larger than 25 nodes are computationally intensive and could require several months to compute (Figure 1C), but, heuristic techniques (stepwise model selection) are promising to drastically speed up the computational time and find a satisfactory solution (Harbord et al., 2016).

The aim of this paper is to provide a fast implementation of DGM, to study its accuracy on a set of simulated networks, and to apply the method on real fMRI data to test reproducibility of the estimated networks. We use Monte Carlo network simulations that have similar properties as real BOLD data to compare network estimates to a ground truth in terms of sensitivity and specificity. We will further compare DGM' performance to another competitive method (Patel's tau; Patel et al., 2006) which was found to demonstrate the highest accuracy for directionality estimation in previous work (Smith et al., 2011). We hypothesize that that DGM will have better detection performance specifically with non-stationary data compared to Patel's tau. We will also perform systematic interventions on the lag of the nodes. We hypothesize that DGM, which are based on instantaneous dependencies between nodes, will not be affected by variation in the haemodynamic response which can confound and reverse directionality estimates, for example as in purely lag-based methods. Having established the validity of DGM, we will estimate dynamic directed relationships of resting-state networks (RSNs) in large samples from the Human Connectome Project to look at the reproducibility of the estimated networks. Finally, we will investigate a hippocampal and a somatosensory network in mouse resting-state fMRI, where we have a more clear expectation of the directionality based on findings from viral tracing studies in the mouse brain.

Materials and methods

Participants and Simulations

We analyzed ten simulated network datasets, each containing time series from a 5-node network of 50 simulated subjects (Table 1). For real fMRI data, we used RSN time series of 500 participants (57% female;

[1] Number of directed graphs with no self-loops, https://oeis.org/A053763.

age distribution of 42%, 36%, 21% and 1% in ranges of 26–30, 31–35, 22–25, and 36+ years) provided by the Human Connectome Project (Smith et al., 2013a; Van Essen et al., 2013). RSN time-series were from the HCP 900 d=25 "PTN" data (Parcellation, Time-series + Netmats), a 25-dimension ICA decomposition of the 15 minutes of resting state data which was acquired with a TR of 0.72 s. From the 25 spatial components we selected ten RSNs which had the highest spatial correlation with the ten RSN's reported by Smith et al. (2009), see Supplementary Table S1. For mouse rs-fMRI data, we analyzed two networks from a previously published dataset of 16 mice, the first is an amygdala-hippocampal-entorhinal network (8 nodes), the second a cortico-striatal-palladium network with 6 nodes (Sethi et al., 2017; Zerbi et al., 2015).

The simulated data follow the same approach as Smith et al. (2011), using a forward model (Friston et al., 2003) that links neural populations to a nonlinear balloon model (Buxton et al., 1998); for full details see Smith et al. (2011). In brief, each node has a binary external input consisting of a Poisson process with an added noise of 1/20 of the neural amplitudes. Mean duration of the simulated firing was 2.5 s (on) and 10 s (off). The neural signals propagate around the network using the DCM with a neural lag of approximately 50 ms. Each node's neural signal was transformed into a simulated BOLD signal using the balloon model resulting in data with approximately 4% signal change, a signal that corresponds to typical fMRI data. These simulations had a hemodynamic response function (HRF) lag variability of 0.5 s, reflecting observed variation of the HR from different brain regions. All simulated data use a 5-node network (Figure 1A). Our first two simulated datasets were "Sim1" and "Sim22" from Smith et al. (2011).[2] Sim1 is a stationary network, and Sim22 an non-stationary (dynamic) network with time-dependent connection strengths. Both network's time series were sampled to create 10 minute time series with a TR of 3 s. In Sim22 the connection strengths are modulated over time by additional random processes, such that the connection strength between two nodes is randomly set to zero for random intervals of mean duration of 30 seconds. For dynamic time series, correlation coefficients between truly connected nodes ranged from 0.25–0.40 across the 5 nodes (mean 0.33), and for stationary time series from 0.30–0.35 (mean 0.32), see Figure 1B.

We created seven additional simulations using the "NetSim" framework of Smith et al. (2011). However, while Smith et al. (2011) randomly changed the delay of each node with a standard deviation of ±0.5 s, we systematically altered the HRF delay at four of the five nodes to create a worst-case, "slow parent, fast child" scenario for a lag-based method. Our interventions systematically increased node 1's delay, and

[2] http://www.fmrib.ox.ac.uk/datasets/netsim/

decreased node 2's delay; the same intervention was performed for node 4 and node 5, and node 3 remained unaffected (Figure 2A). Therefore, child nodes 2 and 5 had their HRF peaks occurring before their respective parent nodes 1 and 4. Node 1 and 4 had increased delays by 0.2–1.2 s compared to the canonical HRF shape, and node 2 and 5 had decreased delays by 0.2–0.6 s (Table 1; Figure 2AC). The seven additional simulated datasets (Table 1) primarily differed regarding the strength of intervention in terms of total systematic offset between connected nodes, with the first dataset having no systematic intervention with an offset of < 0.4 s, and the other simulations having systematic offsets of 0.4 s, 0.8 s, 1.1 s, 1.4 s, 1.7 s, and 1.9 s. Additional descriptive statistics of these datasets are reported in Supplementary Tables S2–S4. The simulated nodes' correlations were comparable to the Sim22 dataset (Figure 2B), and with increased intervention strength, the correlations decreased, as expected. Time series had a mean correlation of r = 0.32 across the five node-pairs and interventions. We also created a long 60 min. simulation of dynamic time series as session duration improved sensitivity in various network modeling methods (Smith et al., 2011).

Dynamic Graphical Models

Dynamic Graphical Models (DGMs) are graphical models with directed relationships and time-varying connectivity weights. Queen and Smith (1993) introduced the MDM, a variant of DGM, which were recently applied to neuroimaging data for the first time (Costa et al., 2015). The time series of a given child node are regressed on the time series from one or more other parent nodes, and the directed relationship corresponds to information flowing from the parent to the child. It is notable for its computational simplicity: Conditional on an easily estimated "discount factor", the model is fully conjugate and model evidence available in closed form.

DGM are specified in terms of a collection of dynamic linear models (DLMs; West and Harrison, 1997):

$$\begin{array}{lll} \text{Observation equations} & Y_t(r) = \boldsymbol{F}_t(r)^T \boldsymbol{\theta}_t(r) + v_t(r) & v_t(r) \sim \mathcal{N}(0, V_t(r)) \\ \text{System equations} & \boldsymbol{\theta}_t(r) = \boldsymbol{\theta}_{t-1}(r) + \boldsymbol{w}_t(r) & \boldsymbol{w}_t(r) \sim \mathcal{N}(0, \boldsymbol{W}_t(r)) \\ \text{Initial information} & (\boldsymbol{\theta}_0(r)|y_0) \sim \mathcal{N}(\boldsymbol{m}_0(r), \boldsymbol{C}_0(r)) & \end{array} \quad (1)$$

where $Y_t(r)$ is the observed value (scalar) of the time series for the child region r at time t, r=1,...,R, t=1,...,T. $\theta_t(r)$ is an unobservable state vector of length |Pa(r)|+1 and represents the connectivity strengths, where

Pa(r) is the set of of parents of node r, and |Pa(r)| is the number of parents . $\mathbf{F}_t(r)$ is a known function of the parents of node r, also of length |Pa(r)|+1; as a regression, $\mathbf{F}_t(r)$ can be thought of as a set of covariates at time t for node r, in our case these are the values of the parent nodes at time t, with the first element of $\mathbf{F}_t(r)$ being 1 to provide an intercept. As a dynamic model, all coefficients $\boldsymbol{\theta}_t(r)$ are allowed to evolve over time, with variability controlled by a scalar "discount factor" δ(r). δ(r) = 1 produces a static model, while for δ(r) = 0.5 the $\theta_t(r)$ evolve as a random walk. A prior distribution which is assumed to be normally distributed with mean $\mathbf{m}_0(r)$ and variance $\mathbf{C}_0(r)$ is assigned for the initial state θ(r); we use the same prior parameter values as in Costa et al. (2015), zero mean and variance of 3 with independence. We conducted prior sensitivity analyses and found these values are suitable for mean zero (at each node) and scaled (globally, over all times and nodes) such that the the average temporal standard deviation (SD) across nodes is 1 (it is crucial not to variance-normalize each node individually, as relative variance contains important directionality information). The observation error $v_t(r)$ and the system error $\mathbf{w}_t(r)$ are normally distributed with variance variance $\mathbf{V}_t(r)$ and $\mathbf{W}_t(r)$ respectively. The error vectors are assumed to be mutually independent, and independent over time and nodes.

The conditional forecast distribution $(Y_t(r)|\mathbf{y}^{t-1}(r), Pa(r))$, that is, the distribution of $Y_t(r)$ given its parents and previous values before time t, $\mathbf{y}^{t-1}(r)$, is a Student *t*-distribution, and the joint log predictive likelihood (LPL) can be calculated as:

$$LPL = \sum_{r=1}^{n} \sum_{t=1}^{T} \log\, p(y_t(r)|\mathbf{y}^{t-1}(r), Pa(r)) \qquad (2)$$

The LPL is the model evidence for a specific set of parents, and as a sum over n nodes, can be optimised for each child node r individually. For each node, for each possible Pa(r), the LPL is first maximised with respect to the discount factor δ(r) using a 1D grid search. While the MDM of Queen and Smith (1993) requires a dynamic acyclic graph (DAG) constraint, finding the best LPL with an acyclic configuration of parents (Pa(1),...,Pa(n)), for the DGM we simply find the maximum LPL for each node individually (i.e. select the highest scoring set of parents). A model m_1 is prefered to m_2 if LPL(m_1) > LPL(m_2). Differences in LPL's correspond to log Bayes factors, and thus provide an interpretable magnitude of evidence for one model over another.

DGM can be regarded as not fitting a single model, but R number of DLMs; this collection of models may not necessarily correspond to a single statistical model (i.e. with positive definite covariance). This pseudolikelihood approach is an approximation to an ideal model, where (e.g.) latent nodes account for reciprocal connections or other sources of cyclic behaviour (Neville and Jensen, 2007). Each fitted DGM, thus, cannot be regarded as single "model", but a set of marginal node models that best describe the directed influence of parents on children nodes.

Pruning networks to improve directionality estimates

While we have found the DGM as defined above to work well on many datasets, we find that it sometimes will estimate bidirectional connections where one direction is much stronger than the other. To simplify interpretation we have considered an optional pruning process where bidirectional connections are converted to unidirectional connection as long as there is only modest reduction in model evidence. The pruning process is as follows. We find the optimal network, i.e. the configuration of parents (Pa(1),...,Pa(R)) that maximises the LPL (2). For each pair of nodes i and j with a bidirectional connection, we compare the model evidence for three possible models *m*:

$$\begin{array}{ll} m_1 : i \leftrightarrows j & LPL(m_1) = LPL_{Pa(i)} + LPL_{Pa(j)} \\ m_2 : i \leftarrow j & LPL(m_2) = LPL_{Pa(i)} + LPL_{Pa(j)\setminus i} \\ m_3 : i \rightarrow j & LPL(m_3) = LPL_{Pa(i)\setminus j} + LPL_{Pa(j)} \end{array} \quad (3)$$

where m_1 is the (LPL-optimal) bidirectional edge, and LPL for these two nodes is the sum of the LPLs of the two winning parent models node *i* and node *j*. m_1 and m_2 are the two models with unidirectional relationships, where each model evidence is likewise the sum of two LPLs but with modified set of parents. For example in m_2, where *i* is the child node, the likelihood is the sum of the LPL of the winning parent model for *i* and the LPL for *j* without *i* as a parent. (Note that the total model evidence for the graph is a sum over all nodes, but since we will only change the model at these two nodes we can disregard the other n-2 contributions to the total LPL for this comparison). The choice between m_1, m_2, and m_3 is determined as follows; first define the evidence of the best unidirectional model as

$$LPL(\hat{m}) = max(LPL(m_2), LPL(m_3)) \quad (4)$$

and then model used is determined as

$$\text{chosen model} = \begin{cases} m_1 & \text{if } LPL(m_1) - LPL(\hat{m}) > e \quad \text{, or} \\ & \text{if } LPL(m_2) = LPL(m_3) \quad \text{, otherwise} \\ m_2 & \text{if } LPL(m_2) > LPL(m_3) \\ m_3 & \text{if } LPL(m_2) < LPL(m_3) \end{cases} \tag{5}$$

where *e* is a chosen penalty, interpreted as the log Bayes Factor comparing the original bidirectional model m_1 to the best unidirectional model $\hat{m}$. If the Bayes Factor for m_1 is more than *e* we retain both the edges, otherwise, we chose a simpler uni-directed model. For the simulation data, we used *e* = 20. For the 10 RSNs in human rs-fMRI, the estimated networks had strong evidence for unidirectional connections, with only 0.6% of edges altered with *e* = 20; as a result we used no pruning (e = 0). The same occurred with the mouse fMRI networks: We used no pruning as *e* = 20 altered only 7% of the edges.

Statistical inference for edge consistency

After estimating the full networks for a dataset, we aim to test the reproducibility of the edges across simulated datasets or over real subjects and runs. We use a binomial test to identify edges that are more prevalent than expected by chance. First, we compute the sample proportion of edges at each directed edge from i to j

$$p_{ij} = \frac{1}{N} \sum_{n=1}^{N} E_{ijn} \tag{6}$$

and define a null edge connection rate as

$$p_0 = \frac{1}{N} \frac{1}{R(R-1)} \sum_{n=1}^{N} \sum_{i \neq j}^{R} E_{ijn} \tag{7}$$

where E_{ijn} is a binary variable indicating the presence/absence of an edge for child node i and parent j from subject n, $1 \leq i \neq j \leq R$, n = 1,...,N. We then test each edge for edge prevalence that differs from the null edge connection rate, with null and alternative hypotheses:

$$H_0 : E(p_{ij}) = p_0 \qquad H_A : E(p_{ij}) \neq p_0 \tag{8}$$

where $E(p_{ij})$ is the true population probability of a specific edge with child *i* and parent *j* across *N* subjects. As we perform multiple testing across edges, we also adjust all *p*-values from the binomial test using false discovery rate (FDR) correction at the 5% level.

Patel's method

We compared DGM's directionality estimates to a method proposed by Patel et al. (2006). Patel's approach is based on pairwise conditional probability: An asymmetry in the probability of A given B, and B given A indicates a directed relationship. As done in previous work (Smith et al, 2011), before calculating these conditional probabilities, the time series were scaled to an interval between 0 and 1 and then binarized with a cutoff at the 0.75 percentile. There are two measures derived from the conditional dependencies, the connection strength (κ) and the direction (τ). We implemented a permutation test for Patel's κ that creates a distribution of κ values under the null hypothesis by shuffling across subjects while keeping the nodes fixed. For each edge with a significant connection strength κ, we assigned the direction based on the sign of τ.

Sensitivity and specificity

We investigated sensitivity and specificity for DGM and Patel's method to detect the true directed network in the simulations. For each simulation setting, we estimated networks from the 50 simulated datasets and measured the number of true positives (TP), false positives (FP), true negatives (TN) and false negatives (FN), where sensitivity is the proportion true positives from condition positive [TP/(TP+FN)], and specificity is the proportion of true negatives from condition negative [TN(TN+FP)].

c-sensitivity and d-accuracy

Smith et al. (2011) defined "c-sensitivity" and "d-accuracy" to compare various methods to detect the true network. We compute and report c-sensitivity, the proportion of correct connections among all simulated datasets, regardless of the directionality. d-accuracy is the fraction of the true positive directionality estimates, this corresponds to what we refer to as "sensitivity".

R package “multdyn”

We released an R (RRID:SCR_001905) package “multdyn” that implements DGM, Patel’s method and all the statistical procedures described above; the package is available on Github and on CRAN[3]. We used version 1.6.1 for all analyses. The number of parent models evaluated for each child node is 2^{n-1}, and hence the total number of evaluations is $n2^{n-1}$; for example a 25-node networks has 419 million possible models. Thus, we implemented some of the time critical functions in C++ using “Rcpp” (Eddelbuettel et al., 2011) and “RcppArmadillo” (Eddelbuettel and Sanderson, 2014). Currently, up to a 20-node network is computationally feasible (20 hours per node on a Intel Xeon CPU (E5-2630 v2) with 2.6 GHz), see Figure 1C. There is a future important extension that we will report separately: we are developing a greedy search algorithm that performs a stepwise search adding or removing a parent node in each step which drastically reduces the search space of all possible parent models (Harbord et al., 2016). The computations can be parallelized on the node level and example jobs are available for both Sun Grid Engine and Slurm on Github. “multdyn” is free software, licensed under GPL3 and runs on all three major platforms (Linux, OS X and Windows).

Results

Network simulations

We compared the performance of the DGM to Patel’s method to detect the true network from the dynamic time series. On dynamic data the DGM (with $e = 20$) had higher sensitivity compared to Patel (DGM: sensitivity 70%, specificity of 79%; Patel: sensitivity 42%, specificity 82%), see Figure 3A. DGM detected the 5 true edges in 14% of the cases, and 4 out of 5 in 56% of the time; Patel’s method detected 0% and 10%, respectively. Looking at individual edges, the highest correct detection rate was 88% for 1→5 and the lowest was 52% for 4→5 for the DGM; Patel’s proportion of correct edges was lower, ranging from 24% (3→4) to 66% (1→5), see Figure 3C. For DGM binomial significance tests on the edge proportions revealed

[3] https://github.com/schw4b/multdyn

all the 5 edges of the true underlying network for dynamic time series. Only the false positive edge 2→1 had a significant proportion of 62%, but was lower compared to the 78% for the true edge 1→2. With Patel's method, the proportion of correct edges was lower, ranging from 24% (3→4) to 66% (1→5), Figure 3C. Only 3/5 true edges had proportions that reached significance, and there was a false positive edge at 1→3. We also evaluated a longer simulation (60 min. instead of 10 min.) and found an increase in sensitivity to 91%, however, specificity was reduced to 60%. When we increased the pruning to reach the same specificity as Patel's method (81%) using $e = 26$, we found a sensitivity of 67%. Using no pruning ($e = 0$), DGM's sensitivity was 84% and specificity was 64% (Figure 3A). For the DGM the overall accuracy was best with $e = 20$ (76%) and $e = 26$ (78%), see Figure 3A. c-sensitivity was higher for DGM (89%) compared to Patel's method (72%). DGM's c-sensitivity for the 60 min. dynamic dataset was 98%.

For stationary data, sensitivity was low for both methods (DGM 50%; Patel: 47%), while Patel's method had higher specificity (DGM: 79%; Patel's method: 86%), see Figure 3B. For DGM, the highest edge proportion was 62% for 2→3 and the lowest was 34% for 3→4, see Figure 3D. Using Patel's method the proportion of correct edges ranged from 60% (1→5) to 30% (3→4). DGM discovered three of the five true edges that reached significant proportions with two significant edges that were false positives. Patel showed 4/5 true edges that reached significance with one false positive edge.

For simulations with the systematic changes to the HRF delay (Figure 4), we found a sensitivity of 80%, 77% and 72% for the first three interventions, node offsets <0.4 s, 0.4 s, 0.8 s, respectively. Stronger interventions reduced sensitivity to 68%, 62% and 55%, for node offset 1.1 s, 1.4 s, 1.7 s, respectively. The strongest intervention with an offset of 1.9 s reduced sensitivity to 48%. Specificity ranged between 62%–69%.

We looked at the discount factors (temporal smoothness) $\delta(r)$ of the connectivity strengths for nodes with at least one parent. Median (IRQ) δ across the 5 nodes was 0.86 (0.84–0.89) for the dynamic data Sim22, and 0.67 (0.66–0.67) for the dynamic simulations with the least intervention (< 0.4 s) and 0.71 (0.65–0.75) for the one with the strongest intervention (1.9 s)

Human resting-state fMRI

The ICA-based resting-state networks from the HCP data had consistent edges in in up to 60% of the participants. The cerebellar, auditory, and also the visual medial networks were child nodes to other RSN's (Figure 5). The visual medial network receives input most prominently from the two other visual networks (visual pole and visual lateral) and the DMN. The visual lateral network is reciprocally connected with the visual medial network and also receives input from the occipital pole. Notably, the estimated networks look very similar across runs. Figure 6 shows the within-subject consistency of the estimated networks, both in terms of presence (Figure 6A) and absence (Figure 6B) of edges. We quantified the edges that occurred in three out of the four runs, and those that occurred in all of the runs. We found that some of the edges described above are reproducible in up 50% of the participants in three runs, and in 30% of the participants in all the four runs. This applied to the reciprocal visual medial and lateral connections and the cerebellum inputs from the DMN, the visual medial, and the auditory network. In terms of reliably absent edges, the visual occipital pole, the DMN, both the right and left frontoparietal network, and also partly the executive control network had consistently no other RSNs as their parents (Figure 6B).

The discount factors $\delta(r)$ of the connectivity strengths for nodes with at least one parent had a median (IQR) across the 10 RSNs of 0.80 (0.74–0.88) in run 1; the runs 2-4 showed similar values.

Mouse resting-state fMRI

We used DGM to investigated dynamic directed functional networks in rs-fMRI data from 16 mice under light anesthesia. We analyzed time series from two networks: the first network has 8 nodes which are regions that communicate with the dentate gyrus (hippocampus). Based on effective connectivity studies, hippocampal areas (DG and CA1) receive information from the entorhinal cortex, the nucleus accumbens and the amygdala, and project to the orbital cortex. Also, DG and CA1 project to the cingulate cortex (Oh et al., 2014). We found a large proportion, 94%, of mice with connectivity from the orbital cortex (OrbM) to the cingulate cortex, see Figure 7A. We also found multiple parent regions of the OrbM area. The CA1 (63% of the mice) and DG (75%) are parent nodes of the cingulate cortex.

The second network contains six nodes that communicate with the motor and somatosensory barrel field areas. Based on effective connectivity studies, the cortical motor and somatosensory areas project to the striatum, and also contralateral flow between the areas can be expected. We found the highest consistency across mice in the contralateral connectivity from the right to the left motor cortex (94%), and from the right to the left putamen (81%), see Figure 7B.

The discount factors δ(r) of the connectivity strengths for nodes with at least one parent had a median (range) across nodes of 0.99 (0.990–0.994) in the 8-node network,; and 0.99 (0.955–0.995) in the 6-node network.

Discussion

We have presented the DGM, a method for modelling dynamic directed functional connectivity. Only a few methods are available to estimate directed relationships. Granger causality for fMRI, as a lag-based method, can be confounded by the variability in the HR of different brain areas, and a latency correction can improve detectability (Wen et al., 2013), but the magnitude of this confound is usually unknown in real data. DCM (Friston, 2009; Friston et al., 2003) implements a biophysical model and can, in principle, infer upon the latent neural mechanisms, but has a high computational effort and is only feasible for a few number of nodes and prior hypotheses. rDCM can estimate larger networks (Frässle et al., 2017) but the effect of HR delay variability has yet to be further investigated. Even though the main focus of this paper is on data from fMRI as it currently is the most widely used method in the neurosciences, DGM are not restricted to this type of data and can potentially be applied in electroencephalography or magnetoencephalography that has a much higher temporal resolution (Baker et al., 2014). In the future, temporal high-resolution data will hopefully capture dynamics that allow a deeper understanding of complex interactions among brain systems, but subsequently require a growing demand for novel methods.

The aim of this study was to investigate the reliability and validity of DGM in simulated and real data and to provide a fast implementation of DGM as an R package. We investigated the performance of the DGM to detect the ground truth network using network simulations. We also compared DGM to a competing method (Patel's method), and we carefully investigated one of the major confounds with directed methods, the temporal variability of the HR. We found that DGM had a sensitivity of 70%–80% and a specificity of

79% and is superior the Patel's method for estimating directionality in the 5-node dynamic simulated networks. DGM detect on average 70% of the edges, and for the edge 1→5 88% of the cases. The edge 4→5 seems to be hardest to detect (52%). It is notable that, from the standpoint of the stationary measure of Pearson's correlation, it was 1→5 that had the highest Pearson correlation (r = 0.40), and 4→5 that had the weakest (r=0.25). The poor performance of the DGM with stationary data can be expected as it is a dynamic method. Interestingly, Patel's tau also performed poorly with stationary data, especially the low sensitivity as compared to the specificity.

Considering the overall accuracy as the only measure of evaluation can be misleading, especially with a sparse network with only 5 out of 20 possible edges. With Patel's method for example, even though detection of edges (sensitivity) was poor, overall accuracy was high due to the high true negative rate. We evaluated sensitivity for both (1) the successful detection of a relationship between nodes and (2) the correct estimation of the directionality simultaneously; while other work (Costa et al., 2015; Smith et al., 2011) looked separately at (1) and (2) and found good directionality estimates (0.60–0.70; d-accuracy) but only under the assumption of (1) being successful. We specifically addressed this issue with a randomization test for Patel's connection strength κ and then assigned the directionality based on the sign of τ.

It is known that variation of the HR across brain regions may confuse lag based methods to estimate directionality. DGM is driven by instantaneous relationships between nodes but does considers some information from the past to regularise state variables (Equation 1; system equation). To test this, we created network simulations with systematic changes on the balloon model priors to increase the lag at parent nodes and decrease the lag at child nodes. We found that if the total offset between a node pair is in the problematic direction (i.e., which would cause Granger to infer incorrect directionality) but below one second, the DGM' sensitivity remains above 70%. If this offset is further increased, sensitivity is declining, but specificity remains intact. Thus, such edges cannot be detected anymore, but incorrectly reversed directionality between the node-pair is not inferred (i.e. the false positive rate).

The major finding for the human resting-state data is that the cerebellar network is most consistently receiving input from many other networks. A meta-analysis demonstrated that the cerebellum is not only relevant in fine motor skills but also involved in cognitive functions, such as working memory, language, emotions, and executive control (Stoodley and Schmahmann, 2009), and subregions of the cerebellum have

been linked to RSNs (Habas et al., 2009). It is important to note that the cerebellar component we used also included appreciable weights for the thalamus and the limbic system, most prominently the amygdala, hippocampus, and the anterior cingulate gyrus. These regions received input from the DMN in a large proportion of subjects, a finding which we also reproduced within-subjects. This suggest a dynamic interplay of memory and emotion associated regions with and triggered by the DMN. Similarly, the cerebellar/limbic network is also triggered by the auditory network. The auditory component involved not only the superior and middle temporal gyrus but also language related left frontal structures and the temporal poles which are highly relevant in semantic memory (Galton et al., 2001). We found that the auditory/language network also receives input from various other networks which could reflect the dynamic access to the language and semantic networks. Furthermore, we found that visual networks have high consistencies in dynamic coupling, such as between the visual medial and the visual lateral network. We found some highly reproducible edges (in up to 94% of the mice), and could successfully determine that CA1 and DG are feeding to the cingulate cortex; this finding is highly consistent with studies that use viral tracers to determine the effective connectivity (Hintiryan et al., 2016; Oh et al., 2014).

In DGM we use a discount factor δ (Petris et al., 2009; West and Harrison, 1997) that sets the temporal smoothness of the connectivity strength θ for each node; this can range from a maximally smooth path (i.e. a constant $\delta(r)\theta = 1$) to a random walk ($\delta(r)\theta = 0.5$; a highly variable path). For each node, for each set of parents, we find the DF that optimizes the LPL, and that optimised LPL is used to identify the optimal set of parents. We found that the δ was higher in Sim22 compared our own dynamic simulations. A possible explanation is that we used a higher sampling rate (lower TR = 2) which preserves more dynamic information in the time series. In the mouse data, we found a δ near 1 suggesting that DGM were not driven by dynamic relationships and was mostly stationary. This may be related to the fact that the mice have been anesthetized (Barttfeld et al., 2015), however, other studies were successful to determine dynamic relationships in anesthetized animals (Hutchison et al., 2013b).

Future studies may investigate dynamic direct connectivity with DGM in task data, for example to classify different tasks based on the directed dynamic coupling of tasks-specific areas or networks. Also, we will be able to estimate larger networks with DGM using a stepwise model search.

Conclusions

Our results demonstrate that DGM discover directed functional connectivity that constitutes reproducible edges in human and mouse rs-fMRI. We observed that in humans the cerebellar/limbic network consistently receives information from other networks, while higher order networks (DMN, parietofrontal) and the V1 are consistently parentless. In network simulations, DGM demonstrated a sensitivity of 72%–77% for dynamic time-series, even in the presence of systemic hemodynamic lag confounds. DGM is a novel and promising approach that constitutes directed relationships providing a more meaningful and mechanistic understanding of brain connectivity in fMRI.

References

Allen, E.A., Damaraju, E., Plis, S.M., Erhardt, E.B., Eichele, T., Calhoun, V.D., 2014. Tracking whole-brain connectivity dynamics in the resting state. Cereb. Cortex 24, 663–676.

Andrews-Hanna, J.R., Smallwood, J., Spreng, R.N., 2014. The default network and self-generated thought: component processes, dynamic control, and clinical relevance. Ann. N. Y. Acad. Sci. 1316, 29–52.

Baker, A.P., Brookes, M.J., Rezek, I.A., Smith, S.M., Behrens, T., Probert Smith, P.J., Woolrich, M., 2014. Fast transient networks in spontaneous human brain activity. Elife 3, e01867.

Barttfeld, P., Uhrig, L., Sitt, J.D., Sigman, M., Jarraya, B., Dehaene, S., 2015. Signature of consciousness in the dynamics of resting-state brain activity. Proc. Natl. Acad. Sci. U. S. A. 112, 887–892.

Bassett, D.S., Bullmore, E.T., 2009. Human brain networks in health and disease. Curr. Opin. Neurol. 22, 340–347.

Bassett, D.S., Gazzaniga, M.S., 2011. Understanding complexity in the human brain. Trends Cogn. Sci. 15, 200–209.

Betzel, R.F., Fukushima, M., He, Y., Zuo, X.-N., Sporns, O., 2016. Dynamic fluctuations coincide with periods of high and low modularity in resting-state functional brain networks. Neuroimage 127, 287–297.

Braun, U., Schäfer, A., Walter, H., Erk, S., Romanczuk-Seiferth, N., Haddad, L., Schweiger, J.I., Grimm, O., Heinz, A., Tost, H., Meyer-Lindenberg, A., Bassett, D.S., 2015. Dynamic reconfiguration of frontal brain networks during executive cognition in humans. Proc. Natl. Acad. Sci. U. S. A. 112, 11678–11683.

Bressler, S.L., Kelso, J.A.S., 2001. Cortical coordination dynamics and cognition. Trends Cogn. Sci. 5, 26–36.

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, 1–38.

Buxton, R.B., Wong, E.C., Frank, L.R., 1998. Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. Magn. Reson. Med. 39, 855–864.

Calhoun, V.D., Miller, R., Pearlson, G., Adalı, T., 2014. The chronnectome: time-varying connectivity networks as the next frontier in fMRI data discovery. Neuron 84, 262–274.

Chase, A., 2014. Altered functional connectivity in preclinical dementia. Nat. Rev. Neurol. 10.

Cheng, W., Rolls, E.T., Qiu, J., Liu, W., Tang, Y., Huang, C.-C., Wang, X., Zhang, J., Lin, W., Zheng, L., Pu, J., Tsai, S.-J., Yang, A.C., Lin, C.-P., Wang, F., Xie, P., Feng, J., 2016. Medial reward and lateral non-reward orbitofrontal cortex circuits change in opposite directions in depression. Brain 139, 3296–3309.

Cocchi, L., Zalesky, A., Fornito, A., Mattingley, J.B., 2013. Dynamic cooperation and competition between brain systems during cognitive control. Trends Cogn. Sci. 17, 493–501.

Costa, L., Smith, J., Nichols, T., Cussens, J., Duff, E.P., Makin, T.R., 2015. Searching Multiregression Dynamic Models of Resting-State fMRI Networks Using Integer Programming. Bayesian Anal. 10, 441–478.

Damaraju, E., Allen, E.A., Belger, A., Ford, J.M., McEwen, S., Mathalon, D.H., Mueller, B.A., Pearlson, G.D., Potkin, S.G., Preda, A., Turner, J.A., Vaidya, J.G., van Erp, T.G., Calhoun, V.D., 2014. Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia. Neuroimage Clin 5, 298–308.

Eddelbuettel, D., François, R., Allaire, J., Chambers, J., Bates, D., Ushey, K., 2011. Rcpp: Seamless R and C++ integration. J. Stat. Softw. 40, 1–18.

Eddelbuettel, D., Sanderson, C., 2014. RcppArmadillo: Accelerating R with high-performance C++ linear algebra. Comput. Stat. Data Anal. 71, 1054–1063.

Fox, M.D., Snyder, A.Z., Vincent, J.L., Corbetta, M., Van Essen, D.C., Raichle, M.E., 2005. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. Proc. Natl. Acad. Sci. U. S. A. 102, 9673–9678.

Frässle, S., Lomakina, E.I., Razi, A., Friston, K.J., Buhmann, J.M., Stephan, K.E., 2017. Regression DCM for fMRI. Neuroimage 155, 406–421.

Friston, K., 2009. Causal modelling and brain connectivity in functional magnetic resonance imaging. PLoS Biol. 7, e33.

Friston, K.J., 2011. Functional and effective connectivity: a review. Brain Connect. 1, 13–36.

Friston, K.J., Harrison, L., Penny, W., 2003. Dynamic causal modelling. Neuroimage 19, 1273–1302.
Friston, K.J., Kahan, J., Biswal, B., Razi, A., 2014a. A DCM for resting state fMRI. Neuroimage 94, 396–407.
Friston, K.J., Stephan, K.E., Montague, R., Dolan, R.J., 2014b. Computational psychiatry: the brain as a phantastic organ. Lancet Psychiatry 1, 148–158.
Galton, C.J., Patterson, K., Graham, K., Lambon-Ralph, M.A., Williams, G., Antoun, N., Sahakian, B.J., Hodges, J.R., 2001. Differing patterns of temporal atrophy in Alzheimer's disease and semantic dementia. Neurology 57, 216–225.
Gonzalez-Castillo, J., Hoy, C.W., Handwerker, D.A., Robinson, M.E., Buchanan, L.C., Saad, Z.S., Bandettini, P.A., 2015. Tracking ongoing cognition in individuals using brief, whole-brain functional connectivity patterns. Proc. Natl. Acad. Sci. U. S. A. 112, 8762–8767.
Guo, C.C., Hyett, M.P., Nguyen, V.T., Parker, G.B., Breakspear, M.J., 2016. Distinct neurobiological signatures of brain connectivity in depression subtypes during natural viewing of emotionally salient films. Psychol. Med. 46, 1535–1545.
Habas, C., Kamdar, N., Nguyen, D., Prater, K., Beckmann, C.F., Menon, V., Greicius, M.D., 2009. Distinct cerebellar contributions to intrinsic connectivity networks. J. Neurosci. 29, 8586–8594.
Handwerker, D.A., Ollinger, J.M., D'Esposito, M., 2004. Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. Neuroimage 21, 1639–1651.
Harbord, R., da Costa, L.C.C., Smith, J., Bijsterbosch, J., Bishop, S., Nichols, T.E., 2016. Scaling up Directed Graph Models for Resting-State fMRI with Stepwise Regression. 18th Annual Meeting of the Organization for Human Brain Mapping.
Hein, G., Morishima, Y., Leiberg, S., Sul, S., Fehr, E., 2016. The brains functional network architecture reveals human motives. Science 351, 1074–1078.
Hintiryan, H., Foster, N.N., Bowman, I., Bay, M., Song, M.Y., Gou, L., Yamashita, S., Bienkowski, M.S., Zingg, B., Zhu, M., Yang, X.W., Shih, J.C., Toga, A.W., Dong, H.-W., 2016. The mouse cortico-striatal projectome. Nat. Neurosci. 19, 1100–1114.
Hutchison, R.M., Womelsdorf, T., Allen, E.A., Bandettini, P.A., Calhoun, V.D., Corbetta, M., Della Penna, S., 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., 2013a. Dynamic functional connectivity: promise, issues, and interpretations. Neuroimage 80, 360–378.
Hutchison, R.M., Womelsdorf, T., Gati, J.S., Everling, S., Menon, R.S., 2013b. Resting-state networks show dynamic functional connectivity in awake humans and anesthetized macaques. Hum. Brain Mapp. 34, 2154–2177.
Huys, Q.J.M., Maia, T.V., Frank, M.J., 2016. Computational psychiatry as a bridge from neuroscience to clinical applications. Nat. Neurosci. 19, 404–413.
Kaiser, R.H., Whitfield-Gabrieli, S., Dillon, D.G., Goer, F., Beltzer, M., Minkel, J., Smoski, M., Dichter, G., Pizzagalli, D.A., 2016. Dynamic Resting-State Functional Connectivity in Major Depression. Neuropsychopharmacology 41, 1822–1830.
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., 2016. On the Stability of BOLD fMRI Correlations. Cereb. Cortex. doi:10.1093/cercor/bhw265
Leonardi, N., Van De Ville, D., 2015. On spurious and real fluctuations of dynamic functional connectivity during rest. Neuroimage 104, 430–436.
Lindquist, M.A., Xu, Y., Nebel, M.B., Caffo, B.S., 2014. Evaluating dynamic bivariate correlations in resting-state fMRI: a comparison study and a new approach. Neuroimage 101, 531–546.
Liu, X., Duyn, J.H., 2013. Time-varying functional network information extracted from brief instances of spontaneous brain activity. Proc. Natl. Acad. Sci. U. S. A. 110, 4392–4397.
Menon, V., 2011. Large-scale brain networks and psychopathology: A unifying triple network model. Trends Cogn. Sci. 15, 483–506.
Nakataki, M., Soravia, L.M., Schwab, S., Horn, H., Dierks, T., Strik, W., Wiest, R., Heinrichs, M., de Quervain, D.J.-F., Federspiel, A., Morishima, Y., 2017. Glucocorticoid Administration Improves Aberrant Fear-Processing Networks in Spider Phobia. Neuropsychopharmacology 42, 485–494.
Neville, J., Jensen, D., 2007. Relational Dependency Networks. J. Mach. Learn. Res. 8, 653–692.
Oh, S.W., Harris, J.A., Ng, L., Winslow, B., Cain, N., Mihalas, S., Wang, Q., Lau, C., Kuan, L., Henry, A.M., Mortrud, M.T., Ouellette, B., Nguyen, T.N., Sorensen, S.A., Slaughterbeck, C.R., Wakeman, W., Li, Y., Feng, D., Ho, A., Nicholas, E., Hirokawa, K.E., Bohn, P., Joines, K.M., Peng, H., Hawrylycz, M.J., Phillips, J.W., Hohmann, J.G., Wohnoutka, P., Gerfen, C.R., Koch, C., Bernard, A., Dang, C., Jones, A.R., Zeng, H., 2014. A mesoscale connectome of the mouse brain. Nature 508, 207–214.

Patel, R.S., Bowman, F.D., Rilling, J.K., 2006. A Bayesian approach to determining connectivity of the human brain. Hum. Brain Mapp. 27, 267–276.

Peterson, A., Thome, J., Frewen, P., Lanius, R.A., 2014. Resting-state neuroimaging studies: A new way of identifying differences and similarities among the anxiety disorders? Can. J. Psychiatry 59, 294–300.

Petris, G., Petrone, S., Campagnoli, P., 2009. Dynamic linear models, in: Dynamic Linear Models with R, Use R. Springer, New York, NY, pp. 31–84.

Preti, M.G., Bolton, T.A., Ville, D.V.D., 2016. The dynamic functional connectome: State-of-the-art and perspectives. Neuroimage. doi:10.1016/j.neuroimage.2016.12.061

Queen, C.M., Smith, J.Q., 1993. Multiregression Dynamic Models. J. R. Stat. Soc. Series B Stat. Methodol. 55, 849–870.

Ramsey, J.D., Hanson, S.J., Hanson, C., Halchenko, Y.O., Poldrack, R.A., Glymour, C., 2010. Six problems for causal inference from fMRI. Neuroimage 49, 1545–1558.

Razi, A., Kahan, J., Rees, G., Friston, K.J., 2015. Construct validation of a DCM for resting state fMRI. Neuroimage 106, 1–14.

Ryali, S., Supekar, K., Chen, T., Menon, V., 2011. Multivariate dynamical systems models for estimating causal interactions in fMRI. Neuroimage 54, 807–823.

Sethi, S.S., Zerbi, V., Wenderoth, N., Fornito, A., Fulcher, B.D., 2017. Structural connectome topology relates to regional BOLD signal dynamics in the mouse brain. Chaos 27, 047405.

Shine, J.M., Bissett, P.G., Bell, P.T., Koyejo, O., Balsters, J.H., Gorgolewski, K.J., Moodie, C.A., Poldrack, R.A., 2016. The Dynamics of Functional Brain Networks: Integrated Network States during Cognitive Task Performance. Neuron 92, 544–554.

Shine, J.M., Poldrack, R.A., 2017. Principles of dynamic network reconfiguration across diverse brain states. Neuroimage. doi:10.1016/j.neuroimage.2017.08.010

Smith, S.M., Beckmann, C.F., Andersson, J., Auerbach, E.J., Bijsterbosch, J., Douaud, G., Duff, E., Feinberg, D.A., Griffanti, L., Harms, M.P., Kelly, M., Laumann, T., Miller, K.L., Moeller, S., Petersen, S., Power, J., Salimi-Khorshidi, G., Snyder, A.Z., Vu, A.T., Woolrich, M.W., Xu, J., Yacoub, E., Uğurbil, K., Van Essen, D.C., Glasser, M.F., WU-Minn HCP Consortium, 2013a. Resting-state fMRI in the Human Connectome Project. Neuroimage 80, 144–168.

Smith, S.M., Fox, P.T., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., Filippini, N., Watkins, K.E., Toro, R., Laird, A.R., Beckmann, C.F., 2009. Correspondence of the brain's functional architecture during activation and rest. Proc. Natl. Acad. Sci. U. S. A. 106, 13040–13045.

Smith, S.M., Miller, K.L., Moeller, S., Xu, J., Auerbach, E.J., Woolrich, M.W., Beckmann, C.F., Jenkinson, M., Andersson, J., Glasser, M.F., Van Essen, D.C., Feinberg, D.A., Yacoub, E.S., Ugurbil, K., 2012a. Temporally-independent functional modes of spontaneous brain activity. Proc. Natl. Acad. Sci. U. S. A. 109, 3131–3136.

Smith, S.M., Miller, K.L., Moeller, S., Xu, J., Auerbach, E.J., Woolrich, M.W., Beckmann, C.F., Jenkinson, M., Andersson, J.L.R., Glasser, M.F., Van Essen, D.C., Feinberg, D.A., Yacoub, E.S., Ugurbil, K., 2012b. Temporally-independent functional modes of spontaneous brain activity. Proc. Natl. Acad. Sci. U. S. A. 109, 3131–3136.

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, 875–891.

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., Uğurbil, K., Van Essen, D.C., 2013b. Functional connectomics from resting-state fMRI. Trends Cogn. Sci. 17, 666–682.

Stephan, K.E., Mathys, C., 2014. Computational approaches to psychiatry. Curr. Opin. Neurobiol. 25, 85–92.

Stephan, K.E., Penny, W.D., Moran, R.J., den Ouden, H.E.M., Daunizeau, J., Friston, K.J., 2010. Ten simple rules for dynamic causal modeling. Neuroimage 49, 3099–3109.

Stoodley, C.J., Schmahmann, J.D., 2009. Functional topography in the human cerebellum: a meta-analysis of neuroimaging studies. Neuroimage 44, 489–501.

Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E.J., Yacoub, E., Ugurbil, K., WU-Minn HCP Consortium, 2013. The WU-Minn Human Connectome Project: an overview. Neuroimage 80, 62–79.

Vidaurre, D., Abeysuriya, R., Becker, R., Quinn, A.J., Alfaro-Almagro, F., Smith, S.M., Woolrich, M.W., 2017. Discovering dynamic brain networks from big data in rest and task. Neuroimage. doi:10.1016/j.neuroimage.2017.06.077

Wen, X., Rangarajan, G., Ding, M., 2013. Is Granger causality a viable technique for analyzing fMRI data? PLoS One 8, e67428.

West, M., Harrison, J., 1997. Bayesian Forecasting and Dynamic Models. Springer New York.

Zalesky, A., Fornito, A., Cocchi, L., Gollo, L.L., Breakspear, M., 2014. Time-resolved resting-state brain networks. Proc. Natl. Acad. Sci. U. S. A. 111, 10341–10346.

Zerbi, V., Grandjean, J., Rudin, M., Wenderoth, N., 2015. Mapping the mouse brain with rs-fMRI: An optimized pipeline for functional network identification. Neuroimage 123, 11–21.

Zhang, J., Cheng, W., Liu, Z., Zhang, K., Lei, X., Yao, Y., Becker, B., Liu, Y., Kendrick, K.M., Lu, G., Feng, J., 2016. Neural, electrophysiological and anatomical basis of brain-network variability and its characteristic changes in mental disorders. Brain 139, 2307–2321.

Zhou, L., Pu, W., Wang, J., Liu, H., Wu, G., Liu, C., Mwansisya, T.E., Tao, H., Chen, X., Huang, X., Lv, D., Xue, Z., Shan, B., Liu, Z., 2016. Inefficient DMN Suppression in Schizophrenia Patients with Impaired Cognitive Function but not Patients with Preserved Cognitive Function. Sci. Rep. 6, 21657.

Figures

Figure 1

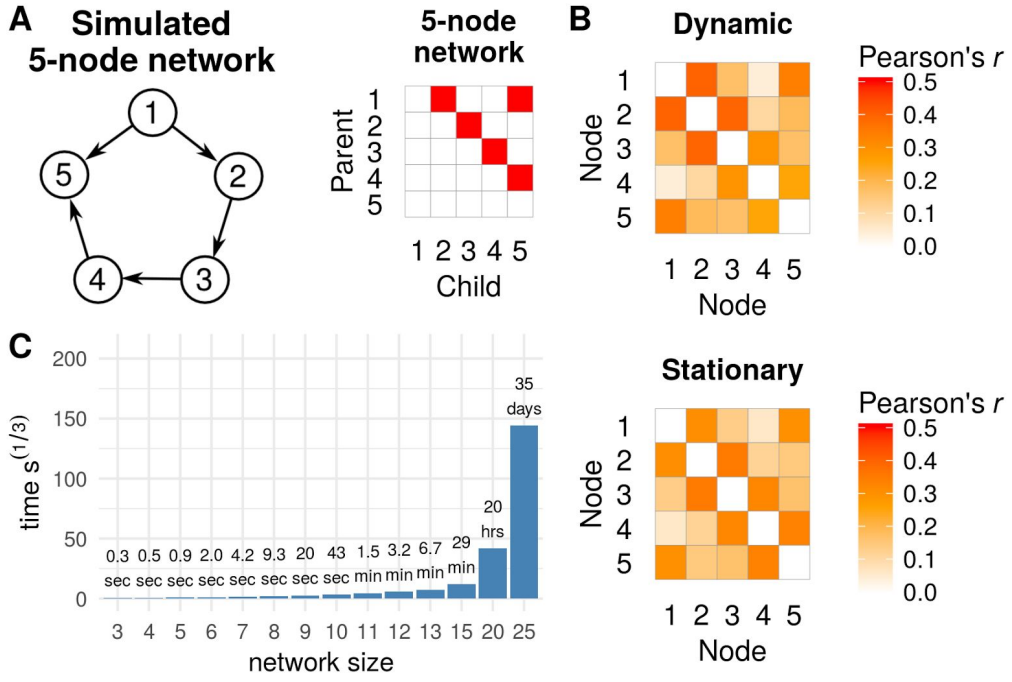

Figure 1. (A) True 5-node network and corresponding adjacency matrix. **(B)** Mean Pearson correlations across 50 simulations for the dynamic (Sim22, top) and stationary network simulations (Sim1, bottom). **(C)** Computational time estimates for a single node of a given network size for a time series of length 1,200 using a Intel Xeon CPU (E5-2630 v2) with 2.6 GHz.

Figure 2

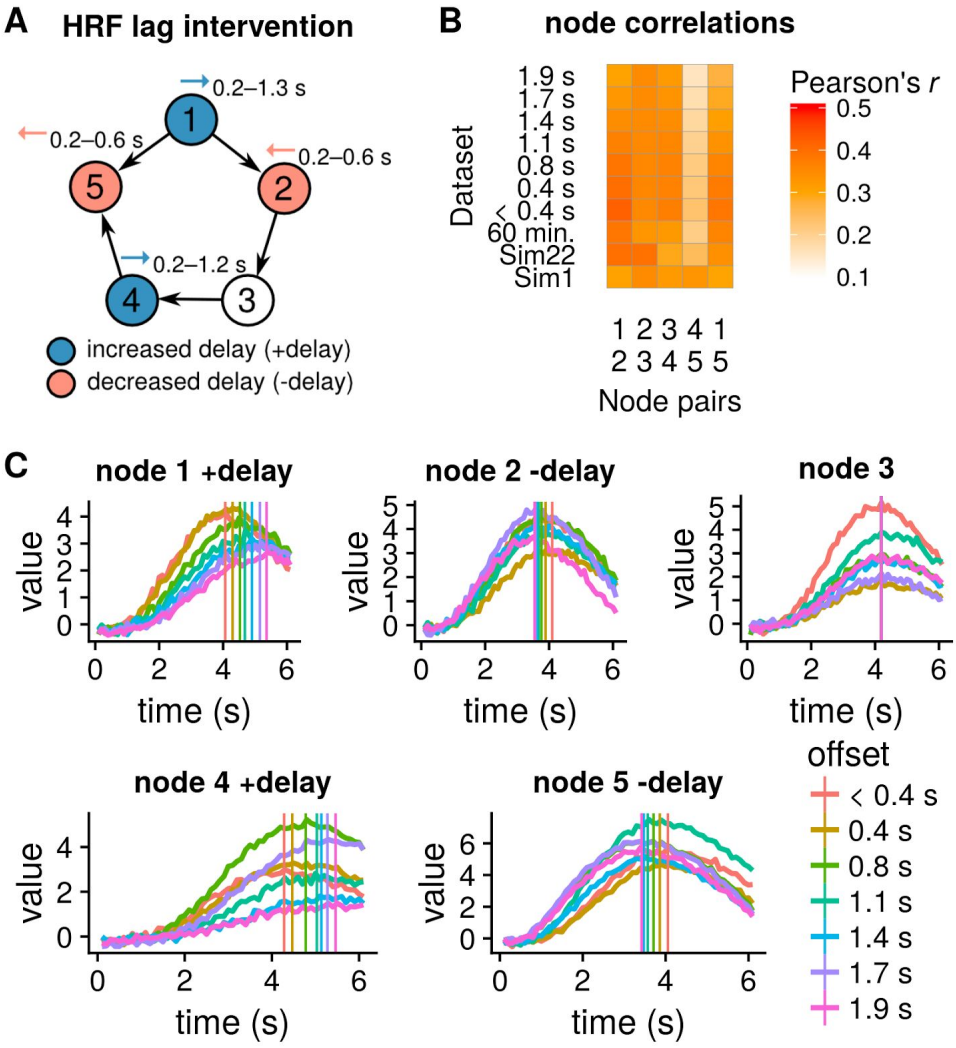

Figure 2. (A) Systematic interventions at four of the five nodes, consisting of either increased (node 1 and 4) or decreased HRF delay (node 2 and 5). We created seven simulations with a total offset between nodes from 0.4–1.9 s **(B)** Pair-wise node correlations in the different simulations. **(C)** HRF shapes of the seven interventions using different DCM forward model priors to shift peaks of the HRF response, with colors from red (< 0.4 s offset) to pink (1.9 s; maximal intervention). Vertical lines indicate average time to peak for each intervention type.

Figure 3

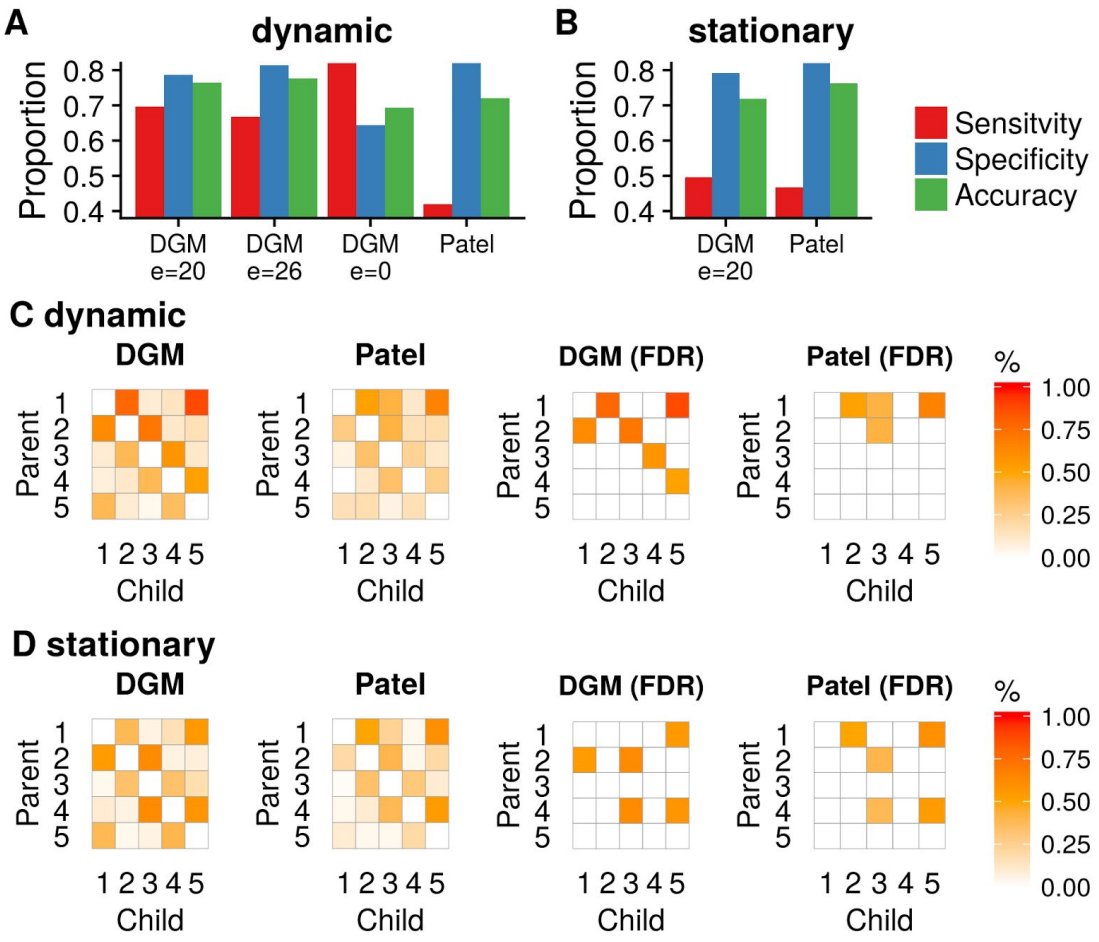

Figure 3. Sensitivity, specificity, and accuracy of the two methods DGM and Patel for dynamic **(A)** and stationary simulations **(B)**. **(C)** Proportions (left) and significant proportions (right, binominal test, 5% FDR threshold) for dynamic network simulations. **(D)** same as **(C)** for stationary network simulations.

Figure 4

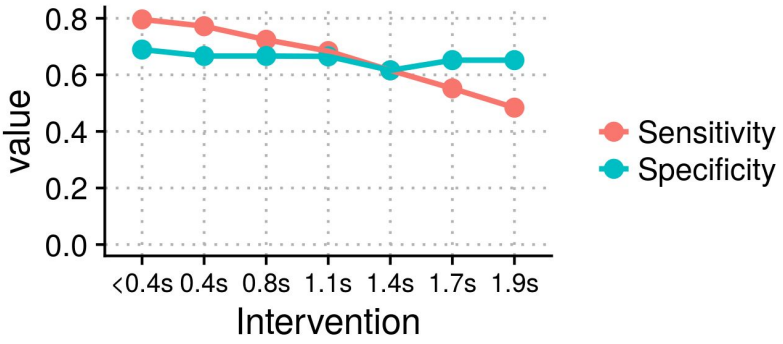

Figure 4. Sensitivity and specificity for DGM (*e* = 20) with different interventions to HRF delays (from 0.4 s up to 1.9 s total offset between two nodes).

Figure 5

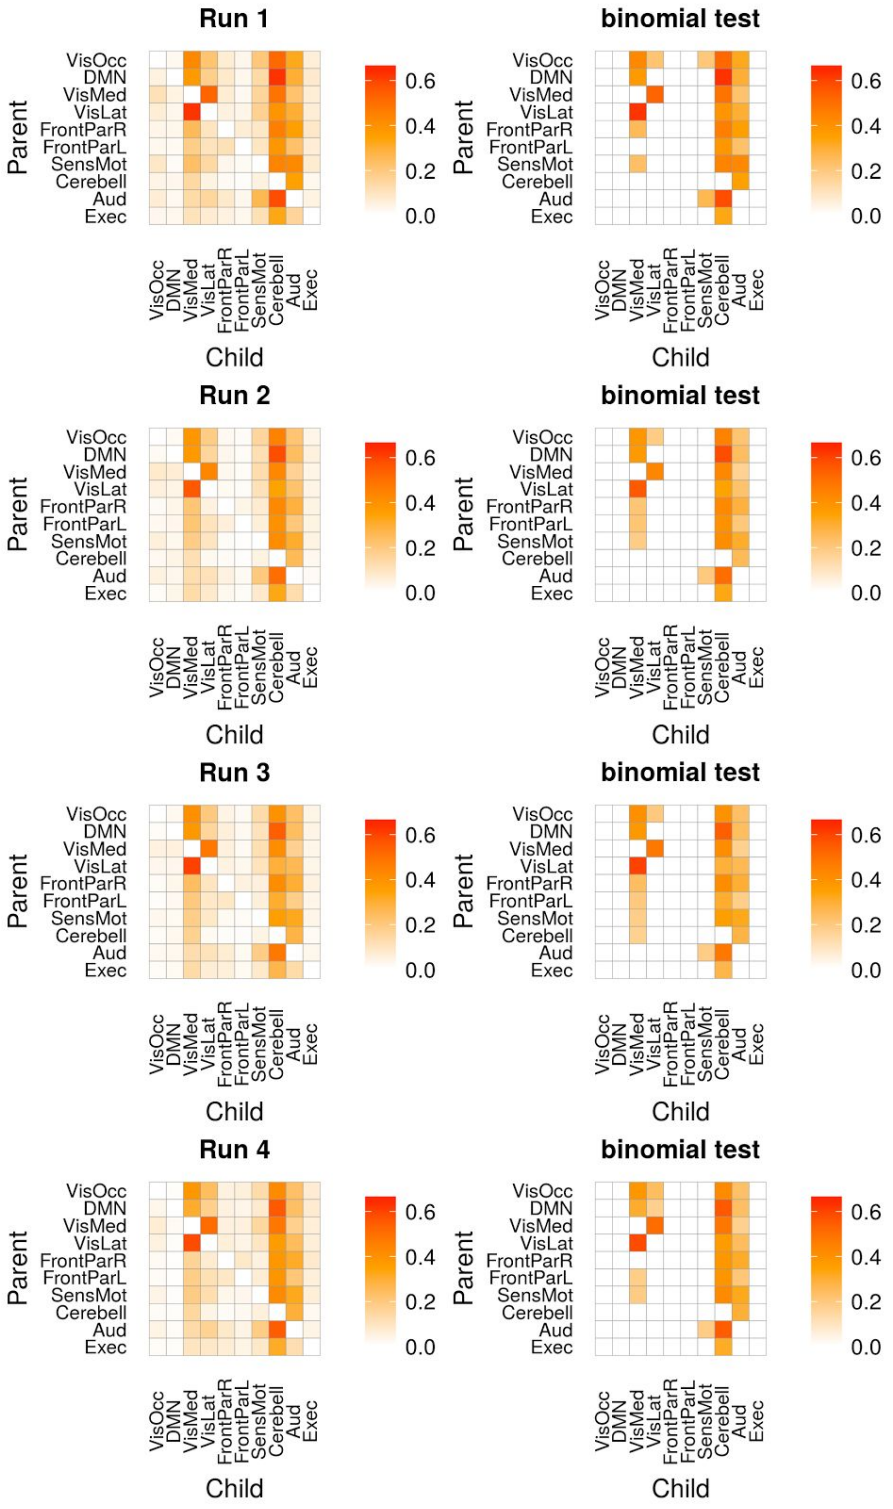

Figure 5. Proportions of edges in participants (*N* = 500) from 4 runs (top to bottom) during resting-state of approximately 15 minutes. Left shows all proportions, right shows proportions for significant edges (binominal test, 5% FDR threshold).

Figure 6

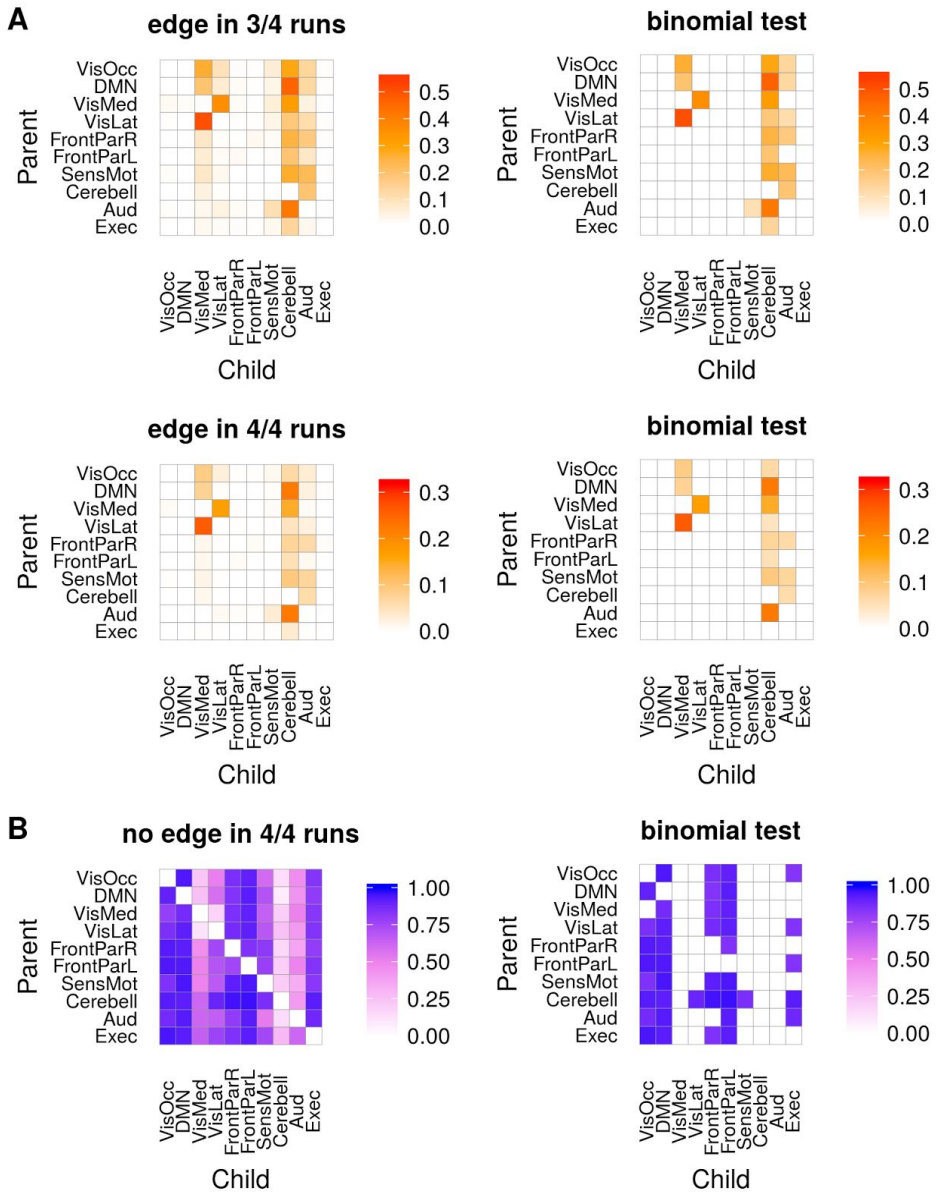

Figure 6. **(A)** Proportions of edges (red) from 500 participants that re-occur in three of the four runs, or in all the four runs (within-subject reproducibility). **(B)** Proportion of participants that had no edge (blue) across all runs. Left shows all proportions, right shows proportions only for significant edges (binominal test, 5% FDR threshold).

Figure 7

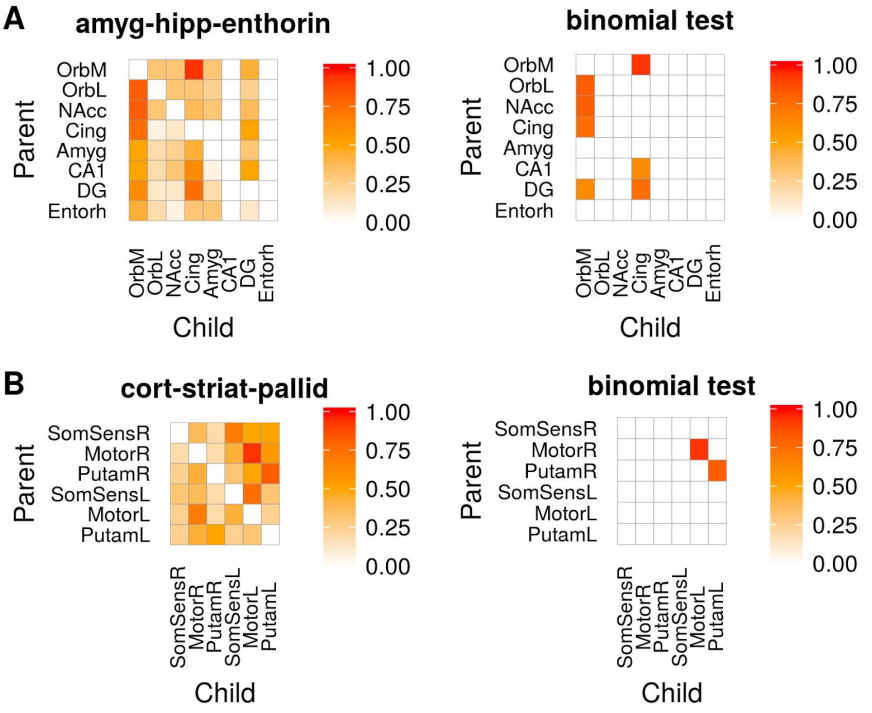

Figure 7. Proportion of edges in fMRI of mice (*N* = 16) testing two networks. **(A)** amygdala-hippocampal-entorhinal network and **(B)** an cortico-striatal network. Left shows all proportions, right shows proportions only for significant edges (binominal test, 5% FDR threshold).

Tables

Table 1

Table 1: Specifications of the simulated BOLD-fMRI data. Each of the simulations contain a simulated 5-node networks for 50 simulated subjects.

Simulation Name	Dynamic network strengths	Duration	TR	Noise	HRF lag variability (s)
Sim1	No	10 min.	3 s	1.0%	±0.5[1]
Sim22	Yes	10 min.	3 s	0.1%	±0.5[1]
< 0.4 s	Yes	10 min.	2 s	0.1%	±0.3[1]
0.4 s	Yes	10 min.	2 s	0.1%	+0.21/-0.20[2]
0.8 s	Yes	10 min.	2 s	0.1%	+0.48/-0.35[2]
1.1 s	Yes	10 min.	2 s	0.1%	+0.69/-0.46[2]
1.4 s	Yes	10 min.	2 s	0.1%	+0.85/-0.54[2]
1.7 s	Yes	10 min.	2 s	0.1%	+1.04/-0.58[2]
1.9 s	Yes	10 min.	2 s	0.1%	+1.24/-0.59[2]
60 min.	Yes	60 min.	2 s	0.1%	±0.3[1]

[1] standard deviation of HRF peak

[2] average increase of node 1 and 4 / average decrease of node 2 and 5