

Directed functional connectivity using dynamic graphical models

Simon Schwab^{1,2,3,*}, Ruth Harbord⁴, Valerio Zerbi⁵, Lloyd Elliott⁶, Soroosh Afyouni¹, Jim Q. Smith², Mark W. Woolrich⁷, Stephen M. Smith⁷, 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 Department of Statistics, University of Warwick, United Kingdom

3 Institute of Digital Healthcare, WMG, 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 and 171598). 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 and analyses can be found at <https://github.com/schw4b/DGM-Sim>. The source code of the R package “DGM” is available at <https://github.com/schw4b/DGM>.

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. This paper proposes dynamic graphical models (DGMs) for dynamic, directed functional connectivity. DGMs are a multivariate graphical model with time-varying coefficients that describe instantaneous directed relationships between nodes. A further benefit of DGMs is that networks may contain loops and that large networks can be estimated. We use network simulations, human resting-state fMRI ($N = 500$) to investigate the validity and reliability of the estimated networks. We simulate systematic lags of the hemodynamic response at different brain regions to investigate how these lags potentially bias directionality estimates. In the presence of such lag confounds (0.4–0.8 seconds offset between connected nodes), our method has a sensitivity of 72%–77% to detect the true direction. Stronger lag confounds have reduced sensitivity, but do not increase false positives (i.e., directionality estimates of the opposite direction). In human resting-state fMRI, we find the DMN has consistent influence on the cerebellar, the limbic and the auditory/temporal network, as well a consistent reciprocal relationship between the visual medial and visual lateral network. Finally, we apply the method in a small mouse fMRI sample and discover a highly plausible relationship between areas in the hippocampus feeding into the cingulate cortex. We provide a computationally efficient implementation of DGM as a free software package for R.

Key words: directed dynamic functional 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 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., 2012a). Emerging evidence suggests that dynamic functional connectivity may reflect changes in the underlying effective connectivity, where changes in effective connectivity cause—or are caused by—crucial aspects of cognition and behaviour (Allen et al., 2014; Braun et al., 2015; Hutchison et al., 2013a; Smith et al., 2012b; Zalesky et al., 2014). The study of these processes can be highly relevant in psychiatry (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., 2017;

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., 2017; 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 FC (Ramsey et al., 2010). Notably, the rate of measurement for fMRI is much slower than the neural time lag, and the hemodynamic response (HR) is subject to variation of the temporal delay. 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 incorrect direction. An interesting issue here is the nature of the lag induced by different haemodynamic response functions in different nodes. Simply changing the parameters or shape of the haemodynamic response function does not necessarily invalidate Granger causality (Barnett and Seth, 2017). However, the introduction of explicit conduction delays violates the temporal precedence assumptions that underlie Granger causality.

Several methods have been proposed to infer causal relationships between nodes, for example exploiting non-Gaussian features of the BOLD signal (Hyvärinen and Smith, 2013; Ramsey et al., 2014) to determine the edge direction between a node pair. The IMAges (Ramsey et al., 2011) method can successfully determine directionality on the group level with concatenated time series across subjects (multi-subject approach), and the GIMME algorithm recovers individual network structure after estimating the group structure (Gates and Molenaar, 2012). A very popular method to estimate directed networks is dynamic causal modeling (DCM; Friston et al., 2003). DCM is primarily designed to make inferences about how experimental conditions change connectivity (Stephan et al., 2010) and is not applicable to resting-state experiments. Usually, only small networks are investigated in DCM studies, as investigators have specific hypotheses on a small number of connections. However, new variants of DCM can be applied to large graphs (rDCM; Frässle et al., 2017; Seghier and Friston, 2013). For example, spectral DCM is suitable to study

networks with dozens of nodes and is designed for resting-state experiments (Friston et al., 2014a; Razi et al., 2017, 2015).

In the present study, we use dynamic graphical models (DGM) to estimate dynamic, directed functional connectivity. There is overwhelming evidence that brain areas are not acyclic due to reciprocal polysynaptic connections (Friston, 2011; Markov et al., 2014). Therefore, unlike our related work on “multiregression dynamic models” (MDMs; Costa et al., 2015a), we do not constrain the networks to be acyclic. Without the acyclic constraint, we refer to our approach as the Dynamic Graphical Model (DGM). A DGM is a dynamic network with time-varying connection strengths that can accommodate larger directed networks. Here, “dynamic” refers to time-varying coupling strengths modeled (as detailed below) by a random walk model of the underlying coupling. In distinction to DCM, we do not consider any time varying experimental conditions or designed input. While this dynamic model is stationary, e.g. random walk smoothness parameters are constant in time, it can explain time-varying correlations in observed fMRI time series. 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. Crucially, 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 $n = 15$ regions/nodes there $2^{n^2-n} \approx 10^{63}$ possible directed graphs¹. 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 $n = 15$) are required. Even so, networks larger than 25 nodes are computationally intensive and could require several months to compute (Figure 1C), but, network discovery (stepwise model selection) are promising to drastically speed up the computational time in the future (Harbord et al., 2016). Also, we show that post-hoc pruning of the networks

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

can further improve directionality estimates; such optimizations are an efficient way to explore model spaces and improve the computational burdens. (Friston et al., 2016; Seghier and Friston, 2013).

The aim of this study is (1) to investigate the accuracy of DGM for recovering the true network structures in simulation based data, (2) to look at the reliability and consistency DGM networks in large samples of real fMRI data, (3) to give a worked example with mouse data and evaluate whether the network structures are plausible and relate to existing literature, and (4) to provide a fast implementation of DGM as an “R” package. Additionally, we include a post-hoc optimization that improves directionality estimates. For the synthetic data, we use generated data from network simulations that have similar properties as real BOLD data (Smith et al., 2011), as we aim to compare network estimates to a ground truth in terms of sensitivity and specificity. We further compare DGM’s performance to other directed methods. We hypothesize that DGM has better detection performance specifically with non-stationary data as the method can model the dynamic relationships between nodes. We also change the simulated node’s hemodynamic lags and hypothesize that the DGM, which is based on instantaneous dependencies between nodes, is not 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 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 investigate a hippocampal and a somatosensory network in mouse resting-state fMRI, where we have a clear expectation of the directionality based on findings from viral tracing studies in the mouse brain.

Materials and methods

Network simulations

We created various synthetic fMRI datasets based on network simulations of a 5-node ($N = 50$ datasets) and a 8-node network ($N = 20$), see Table 1. 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 2.7 s (on) and 10 s (off); with a range between 770 ms and 10.1 s. The network also included self-connections to model within-node temporal decay; this was set to approximately 50 ms for the 5-node network, and 500 ms for the 8-node network. These specifications of neural activity is rather ad hoc, but these capture representative settings suitable for our evaluations. 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 lag variability of 0.5 s, reflecting observed variation of the HR from different brain regions. The first two simulated 5-node networks (Figure 1A) 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 with systematic offsets between the nodes using the “NetSim” framework of Smith et al. (2011). With “offset”, we refer to the difference between two nodes with respect to their hemodynamic delays. For example, if we increase the HR delay at node 1 by +0.5 s and decrease the delay at node 2 by -0.2 s, the resulting total offset introduced between the two nodes is 0.7 s. The HR delay is the increased/decreased lag compared to the canonical HRF at a node, which is usually around 4–6 s. 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”

² <http://www.fmrib.ox.ac.uk/datasets/netsim/>

scenario for a lag-based method (Figure 2A). To achieve this, we changed the shape of the hemodynamic response function: we changed the transit time of blood τ in the balloon model, i.e., the average time blood takes to traverse the venous compartment (Buxton et al., 1998; Stephan et al., 2007). Values for τ were sampled from a normal distributions with means provided in the Supplementary Table S2 and a variance of 0.05. For node 1 and 4, we increased τ , and for node 2 and 5 we decreased τ . We confirmed this by a simulation using a single block of neural spiking with a duration of 2.5 s and measured the time to peak in the simulated HR at each node as the signal moves through the network; see Figure 2C and Supplementary Table S3. As a result our interventions increased node 1's delay, and decreased node 2's delay; the same intervention was performed for node 4 and node 5, and node 3 remained unaffected (Figure 2AC). Therefore, nodes 2 and 5 had their HRF peaks occurring before their respective parent nodes 1 and 4. In our simulations that included these offsets, node 1 and 4 had increased delays by 0.21–1.24 s compared to the canonical HRF shape, and node 2 and 5 had decreased delays by 0.20–0.59 s. These 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 S3–S5. 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).

Finally, we extended the 5-node network and created a stationary and nonstationary 8-node network simulation with two reciprocal connections, and a neural lag of 500 s (Figure 7A). We aimed to investigate whether DGM can successfully detect the bidirectional relationships. The neural lag was changed from 50 ms to 500 ms as this more reasonably reflects neural activity by populations of neurons with slower fluctuations.

fMRI data

For real fMRI data, we used resting state network (RSN) time series of 500 participants (57% female; 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 RSNs reported by Smith et al. (2009), see Supplementary Table S1. Note that we are treating a RSN as a node in our analyses, i.e., we are summarising the activity of a distributed mode (i.e., resting state network) with a single time series and then using the DGM to understand how 'networks' are coupled to 'networks'. The reason we can do this is that most resting state networks comprise spatially compact regions that can be treated as a small collection of nodes. 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-pallidum network with 6 nodes (Sethi et al., 2017; Zerbi et al., 2015).

Dynamic Graphical Models

Dynamic Graphical Models (DGM) are graphical models with directed relationships and time-varying connectivity weights. The connectivity weights are the regression coefficients that reflect the effect of a set of parent nodes as covariates on a child node; see Figure 3. DGM are closely related to the Multiregression Dynamic Model (Costa et al., 2015b; MDM; Queen and Smith, 1993), except that a DGM need not be a directed acyclic graph (DAG). A DGM is a set of dynamic linear models (DLMs), state space models that are linear and Gaussian (Petrus et al., 2009a; West and Harrison, 1997). In a DLM, the time series of a given child node are regressed on the time series from one or more other parent nodes using dynamic regression. The directed relationship corresponds to information flowing from the parent to the child. A DGM is notable

for its computational tractability: conditional on an easily estimated ‘discount factor’, the model is fully conjugate and the model evidence available in closed form.

We define a DGM as a collection of multiple regression DLMS described by the following set of equations (Harrison and Stevens, 1976; West and Harrison, 1997; p. 109):

$$\begin{aligned}
 \text{Observation equation} \quad & Y_t(r) = \mathbf{F}_t(r)^T \boldsymbol{\theta}_t(r) + v_t(r) \quad v_t(r) \sim \mathcal{N}(0, V(r)) \\
 \text{System equation} \quad & \boldsymbol{\theta}_t(r) = \boldsymbol{\theta}_{t-1}(r) + \mathbf{w}_t(r) \quad \mathbf{w}_t(r) \sim \mathcal{N}(\mathbf{0}, V(r) \mathbf{W}_t^*(r)) \\
 \text{Initial information} \quad & (\boldsymbol{\theta}_0(r) | y_0, V(r)) \sim \mathcal{N}(\mathbf{m}_0(r), V(r) \mathbf{C}_0^*(r)) \\
 \text{Prior on the precision} \quad & (\phi(r) | y_0) \sim \mathcal{G}\left(\frac{n_0(r)}{2}, \frac{d_0(r)}{2}\right)
 \end{aligned} \tag{1}$$

For some child region r , $Y_t(r)$ is the observed value at time t , ($r = 1, \dots, n$, $t = 1, \dots, T$). Let $Pa(r)$ denote the parents of r and $|Pa(r)|$ be the number of parents. Then $\boldsymbol{\theta}_t(r)$ is an unobservable system vector with length $|Pa(r)| + 1$. The covariate vector $\mathbf{F}_t(r)$ has the same dimension as $\boldsymbol{\theta}_t(r)$ and contains the observations from the parent nodes $Pa(r)$ at time t ; the first element of $\mathbf{F}_t(r)$ is 1 to provide an intercept. The observation error $v_t(r)$ and the system error $\mathbf{w}_t(r)$ are normally-distributed with variances $V(r)$ and $\mathbf{W}_t(r) = V(r) \mathbf{W}_t^*(r)$ respectively. The error vectors are assumed to be mutually independent and independent over time and nodes. The observation variance $V(r)$ is unknown and assumed to be constant so that we may place a gamma-distributed prior on the precision $\phi(r) = V(r)^{-1}$ and express all of the variances in the DLM in terms of this constant, unitless term (West and Harrison, 1997). Finally, the initial system vector $\boldsymbol{\theta}_0(r)$, conditioned on y_0 , is the available information at time $t = 0$, and represents prior knowledge of the regression coefficients before observing any data; $\boldsymbol{\theta}_0(r)$ follows a normal distribution with mean $\mathbf{m}_0(r)$ and variance $\mathbf{C}_0 = V(r) \mathbf{C}_0^*(r)$. The system variance is re-parameterized in terms of the posterior variance $\mathbf{C}_{t-1}(r) = V(r) \mathbf{C}_{t-1}^*(r)$ of the state variables $\boldsymbol{\theta}_t(r)$ at time $t - 1$

$$\mathbf{W}_t^*(r) = \frac{1 - \delta(r)}{\delta(r)} \mathbf{C}_{t-1}^*(r) \tag{2}$$

where $\delta(r)$ is a scalar ‘discount factor’. From this expression, it can be seen that $\delta(r) = 1$ corresponds to a static model, while lower values allow dynamic connectivity estimates.

The prior hyperparameters $\mathbf{m}_0(r)$, $\mathbf{C}_0^*(r)$, $n_0(r)$, and $d_0(r)$ must be specified *a priori*. We used the same values as Costa et al. (2015b) which are $\mathbf{m}_0(r) = \mathbf{0}$, $\mathbf{C}_0^*(r) = 3I_{|Pa(r)|+1}$ where I is the identity matrix, $n_0(r) = 0.001$ and $d_0(r) = 0.001$. We conducted prior sensitivity analyses and found these values are suitable for mean zero time series at each node and scaled (globally, over all time series and nodes) such that 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 forecast distribution for $Y_t(r)$ given parents $Pa(r)$ and past observations $\mathbf{y}^{t-1}(r)$ follows a Student’s t-distribution (West and Harrison, 1997). This distribution allows us to evaluate the log likelihood of an observed time series by summing over time and regions. For an individual node r with parents $Pa(r)$, the log likelihood, ie. the model evidence (ME), is

$$\text{ME}(r) = p(y_1(r), y_2(r), \dots, y_T(r) | Pa(r)) = \sum_{t=1}^T \log p(y_t(r) | \mathbf{y}^{t-1}(r), Pa(r)) \quad (3)$$

so that the model evidence for some model m with parents $Pa = (Pa(1), \dots, Pa(n))$ is expressed as follows:

$$\text{ME} = \sum_{r=1}^n \text{ME}(r) \quad (4)$$

For each possible $Pa(r)$, the ME is first maximised with respect to the discount factor $\delta(r)$ using a 1D grid search, $\delta(r) \in [0.5, 1]$. While the MDM of Queen and Smith (1993) requires a DAG constraint, finding the highest ME with an acyclic configuration of parents ($Pa(1), \dots, Pa(n)$), for DGM we simply find the maximum ME for each node individually (i.e. select the highest scoring set of parents). A model m_1 is preferred to m_2 if $\text{ME}(m_1) > \text{ME}(m_2)$. Differences in ME correspond to log Bayes factors, and thus provide an interpretable magnitude of evidence for one model over another. In other words, a model is defined by the parents of a node; namely, whether there can be an influence or not. Given the model, it is possible to optimise the parameters in Equation 1, and evaluate the model evidence in Equation 3 and 4. The differences

in model evidence can be used to directly compare models in terms of log Bayes factors (assuming that they have the same prior probability).

As DGM is fitting a collection of models, one per node, this collection may not necessarily correspond to a single statistical model (i.e. with a positive definite covariance posterior at every time point). This pseudolikelihood approach (Neville and Jensen, 2007) is an approximation to an ideal model, where (e.g.) latent nodes are introduced to account for reciprocal connections or other sources of cyclic behaviour. Each fitted DGM, thus, is best interpreted as a set of marginal node models that best describe the directed influence of parents on children nodes. This factorization of DGM model evidence by node is the same device used in regression DCM (Frässle et al., 2017).

Example of a DGM with 3 nodes

Let's assume a 3-node network with observed time series at the three nodes $Y_t(1)$, $Y_t(2)$, and $Y_t(3)$, see Figure 3. For example, node 3 has node 1 and node 2 as parents. Given the graphical structure, the model equations for the three nodes are written in Figure 3. The time-varying coefficients of the three edges $e_{1 \rightarrow 2}$, $e_{1 \rightarrow 3}$, $e_{2 \rightarrow 3}$ are $\theta_t^{(2)}(2)$, $\theta_t^{(2)}(3)$, and $\theta_t^{(3)}(3)$, respectively. These estimates are for every time t , allowing the connectivity strengths to evolve over time.

Model-comparison based pruning of reciprocal connections

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. We find the optimal network, i.e. the configuration of parents ($Pa(1), \dots, Pa(n)$) that maximises the model evidence (ME). For each pair of nodes i and j with a bidirectional connection, we compare the model evidence for three possible models m (see also Figure 4):

$$\begin{aligned} m_1 : i &\leftrightarrow j & ME(m_1) &= ME_{Pa(i)} + ME_{Pa(j)} \\ m_2 : i &\leftarrow j & ME(m_2) &= ME_{Pa(i)} + ME_{Pa(j) \setminus i} \\ m_3 : i &\rightarrow j & ME(m_3) &= ME_{Pa(i) \setminus j} + ME_{Pa(j)} \end{aligned} \tag{5}$$

where m_1 is the (ME-optimal) bidirectional edge, and the ME for these two nodes is the sum of the MEs 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 MEs but with modified set of parents. For example in m_2 , where i is the child node, the likelihood (model evidence) is the sum of the ME of the winning parent model for i and the ME 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 ME for this comparison. The choice between m_1 , m_2 , and m_3 is determined as follows: first, we define the evidence of the best unidirectional model as

$$ME(\hat{m}) = \max(ME(m_2), ME(m_3)) \quad (6)$$

and then the model used is determined as

$$\text{chosen model} = \begin{cases} m_1 & \text{if } ME(m_1) - ME(\hat{m}) > e, \text{ or} \\ & \text{if } ME(m_2) = ME(m_3) \\ m_2 & \text{if } ME(m_2) > ME(m_3) \\ m_3 & \text{if } ME(m_2) < ME(m_3) \end{cases} \quad (7)$$

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. One can regard e as a prior over unidirectional versus bidirectional connections. In other words, if we use the value of $e = 20$, we are effectively saying that our prior odds for unidirectional connections over bidirectional connectivity is $\exp(20)$. For the simulation data, we used $e = 20$ for the 5-node network, and $e = 10$ for the 8-node network. 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

$$p_{ij} = \frac{1}{N} \sum_{n=1}^N E_{ijn} \quad (8)$$

and define a null edge connection rate as

$$p_0 = \frac{1}{N} \frac{1}{R(R-1)} \sum_{n=1}^N \sum_{i \neq j}^R E_{ijn} \quad (9)$$

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

$$H_0 : E(p_{ij}) = p_0 \quad H_A : E(p_{ij}) \neq p_0 \quad (10)$$

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.

Comparison with other directed methods

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 node A given node B, and B given A indicates a directed relationship. The model for the joint activation of each node pair is based on a multinomial likelihood with a Dirichlet prior distribution. 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 the time series 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 .

We also compared DGM to the pairwise Lingam method (Hyvärinen and Smith, 2013), an extension to the standard structural equation model that exploits non-Gaussianity of the data.

Sensitivity and specificity

We investigated sensitivity and specificity for DGM and the various other network estimation methods to detect the true directed network in the simulations. For each simulation setting, we estimated networks from the 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 $[TN/(TN+FP)]$. Additionally, we computed “c-sensitivity” (Smith et al., 2011), the proportion of correct connections among all simulated datasets, regardless of the directionality.

R package “DGM”

We released an R (RRID:SCR_001905) package “DGM” that implements DGM, Patel’s method and all the statistical procedures described above; the package is available on Github and on CRAN³. 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 have developed a greedy search algorithm that performs a stepwise search adding or removing a parent node in each step which drastically

³ <https://github.com/schw4b/DGM>

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. “DGM” is free software, licensed under GPL3 and runs on all three major platforms (Linux, OS X and Windows).

Results

Network simulations

We evaluated the performance of the DGM to detect the true relationships of the non-stationary and stationary network nodes and compared to Patel’s and the Lingam method. On dynamic data the DGM (with $e = 20$) had higher sensitivity compared to Patel and Lingam (DGM: sensitivity and specificity, 70% and 79%; Patel: 42% and 82%; Lingam: 41% and 80%), see Figure 5A. 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%, and Lingam 0% and 8%, respectively. Looking at individual edges, the highest correct detection rate was 88% for $1 \rightarrow 5$ and the lowest was 52% for $4 \rightarrow 5$ for the DGM; Patel’s proportion of correct edges was lower, ranging from 24% ($3 \rightarrow 4$) to 66% ($1 \rightarrow 5$), and Lingam from 26% ($1 \rightarrow 2$) to 60% ($4 \rightarrow 5$), see Figure 5C. For DGM binomial significance tests on the edge proportions revealed all the 5 edges of the true underlying network for dynamic time series. Only the false positive edge $2 \rightarrow 1$ had a significant proportion of 62%, but was lower compared to the 78% for the true edge $1 \rightarrow 2$. With Patel’s method, only 3/5 of the true edges had proportions that reached significance, and there was a false positive edge at $1 \rightarrow 3$; Lingam discovered 3/5 edges but had 5 false positives and mainly estimated the inverse directions. When we increased the pruning to reach the same specificity as Patel’s method (82%) using $e = 26$, we found a sensitivity of 67% for DGM. Using no pruning ($e = 0$), DGM’s sensitivity was 84% and specificity was 64% (Figure 5A). For the DGM, the overall accuracy was best with $e = 20$ (76%) and $e = 26$ (78%), see Figure 5A. c-sensitivity was higher for DGM (89%) compared to Patel’s method (72%); we did not evaluate c-sensitivity for Lingam as this method only

determined directionality. We also evaluated a longer simulation with DGM (60 min. instead of 10 min.) and found an increase in sensitivity to 91%, however, specificity was reduced to 60%; c-sensitivity was 98%.

For stationary data, performance was best for Lingam with a sensitivity of 70% (DGM 50%; Patel: 47%), and a specificity of 90% (DGM: 79%; Patel: 86%), see Figure 5B. For DGM, the highest edge proportion was 62% for 2→3 and the lowest was 34% for 3→4, see Figure 5D. Using Patel's method the proportion of correct edges ranged from 30% (3→4) to 60% (1→5), and for Lingam from 58% (4→5) to 84% (1→2). DGM discovered three of the five true edges that reached significant proportions with two significant edges that were false positive Patel showed 4/5 true edges that reached significance with one false positive edge, and Lingam recovered all true edges and one false positive.

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 $\delta(r)$, the temporal smoothness of the connectivity strengths $\theta(r)$ for nodes with at least one parent. Median $\delta(r)$ across the 5 nodes was 0.86 (0.84–0.89) (Q1–Q3) 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). For the stationary simulation (Sim1), a median $\delta(r)$ of 0.994 (0.992–0.996) was fitted.

We also tested a 8-node network with reciprocal connections (Figure 7A). The reciprocal connections showed higher correlation compared to the unidirectional edges, which was expected (Figure 7B). On dynamic data the DGM (with $c = 10$) had higher sensitivity compared to Lingam (DGM: sensitivity and specificity, 81% and 73%; Lingam: 49% and 93%), see Figure 5E. However, note here that Lingam cannot detect a relationship between two nodes, only the direction of the edge, and therefore we informed Lingam with the edges that have a true connection. For DGM binomial significance tests on the edge proportions revealed all the 10 edges of the true underlying network for dynamic time series with 3 false positives. With

Lingam, only 7/10 of the true edges had proportions that reached significance, and there were 5 false positive edges. On stationary data the DGM (with $e = 10$) had higher sensitivity compared to Lingam (DGM: sensitivity and specificity, 80% and 65%; Lingam: 65% and 97%), see Figure 5E. After significance testing, DGM revealed all the 5/10 edges with 6 false positives, and Lingam recovered all the 10 edges that are true with 2 false positives, see Figure 7D.

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 8). 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 9 shows the within-subject consistency of the estimated networks, both in terms of presence (Figure 9A) and absence (Figure 9B) 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 9B).

The discount factors $\delta(r)$ of the connectivity strengths for nodes with at least one parent had a median (Q1-Q3) 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 investigate 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 are functionally connected with the dentate gyrus of the hippocampus (DG). Directed anatomical connectivity based on tracer data suggests that hippocampal areas (DG and CA1) receive projections from the entorhinal and other associative cortices (Weilbacher and Gluth, 2016), the nucleus accumbens and the amygdala, and project to prefrontal areas (Eichenbaum, 2017; Hintiryan et al., 2016; Jin and Maren, 2015; Oh et al., 2014). We found that the DG received inputs from several nodes, but missed significance. However, a large proportion, 94%, of mice showed directed connectivity from the orbital cortex (OrbM) to the cingulate cortex, see Figure 11A. We also found multiple parent regions of the OrbM area, including the DG (63% of mice). Furthermore, the node CA1 (63% of the mice) and DG (75%) were parent nodes of the cingulate cortex.

The second network contains six nodes that communicate with the motor and somatosensory barrel field areas. Based on axonal/white-matter connections, the cortical motor and somatosensory areas project to the striatum, and also contralateral flow between the areas can be expected (Lee et al., 2016). 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 11B.

The discount factors $\delta(r)$ 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

Dynamic graphical models (DGM) are a collection of dynamic linear models (DLM) at each node where we treat parent nodes as covariates in a dynamic multiple regression. Even though the (graphical) model structure is static once we optimized the family of parental nodes, we allow for time-varying correlations in

the observed fMRI time-series. We extended previous work by (Costa et al., 2015a), relaxing the DAG constraint and allowing for cycles. Furthermore, we introduced systematic lags to test robustness against hemodynamic confounds and we could successfully invert time-varying node relationships in simulation data generated by a the Smith et al. (2011) generative model based on a deterministic DCM. In determining directed relationships we were partly successful, but this remains a challenging task. We could establish consistent relationships across a large number of subjects that demonstrated top-down influence from the DMN to the cerebellum and limbic regions, as well as to temporal and memory related areas; and we found a reciprocal connection of the visual medial and visual lateral areas. In the future, DGM has the potential to estimate large networks with 100 of nodes (by optimization of network discovery) which can be suitable for further analyses with graph theory (Bullmore and Sporns, 2009; Fornito et al., 2016). Hopefully, this device allows to study pathological network topologies of brain disorders in the future (Fornito et al., 2015), but further validation of the methods are required.

We found that DGM had a sensitivity and specificity between 70%–80% for the 5 and 8-node network with nonstationary node relationships. Patel's method suffered low sensitivity to detect a relationship between nodes but can establish the correct directionality in a majority of edges. Lingam oriented the network correctly, but only in the stationary case. However, one limitation here is that Lingam is not a full model, i.e. the approach requires to know the edges a priori and then orients those based on non-gaussian properties.

DGM showed lower sensitivity with stationary data in the 5-node network but not for the 8-node network. This is interesting as DGM has the ability to fit a fully stationary model, and indeed such a stationary model was generally fitted with the stationary data. However, the directionality structure then often suffers in this case and it appears the dynamic information in the data is highly relevant in the determination of directionality.

The main 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

functionally related 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, a finding which was consistent across subjects and also within-subject across the four runs. This could reflect a intrinsic dynamic coupling of regions associated with memory and emotion with the DMN (Sestieri et al., 2011). The auditory component not only involved 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 used DGM to explore directed functional connectivity in the mouse brain. 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 consistent with studies that use viral tracers to determine the directed anatomical connectivity (Hintiryan et al., 2016; Oh et al., 2014). A limitation here is that we compare functional connectivity with findings from anatomical connectivity which are conceptually different and certainly have not a one-to-one relationship, however, there is evidence from functional studies that the hippocampus and the anterior cingulate cortex are prominent areas of the limbic circuit and the DMN.

In DGM we use a discount factor δ (Petrus et al., 2009b; 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 $\delta(r) = 1$) to a highly variable path ($\delta(r) = 0.5$). For each node, and for each set of parents, we find the DF that optimizes the model evidence (ME), and that optimised ME 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).

Various methods exist to establish directed relationships: Granger causality for fMRI, as a lag-based method, can be confounded by the variability in the HR of different brain areas. 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. However, recent developments allow DCM to study larger networks and also resting-state fMRI (Frässle et al., 2017; Razi et al., 2017). In future work, we would like to establish construct validity between DGM and spectral DCM (spDCM) for resting state (Friston et al., 2014a; Razi et al., 2015). In the current paper, we focused on comparing DGM to Patel's method and Lingam within the scope of purely observational models of directed functional connectivity. In contrast, spDCM, which aims to infer effective connectivity, has a different agenda which is to infer causal influences between different neuronal populations. spDCM is a hierarchical model with two sets of neural and observation (hemodynamic) equations. spDCM assumes neural fluctuations to have a power law (scale-free) distribution whereas the deterministic (forward) DCM for our simulations had an exogenous input. Therefore, our current simulations will not be compatible with the generative model employed in spDCM. Another method to determine directed functional connectivity is Group Iterative Multiple Model Estimation (GIMME), a structural equation model (Gates and Molenaar, 2012). Here, information on an estimated group-level network is used to improve network structure estimates on the subject-level. The method showed a high sensitivity and accuracy (both above 80%), but we have not evaluated it here as it does not provide subject-level estimates.

Considering the overall accuracy as the only measure of evaluation can be biased, 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., 2015a; 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 .

A particular problem to fMRI is the variation of the hemodynamic response across brain regions that may a confounding factor for lag based methods to estimate directionality. DGM, however, is driven by instantaneous relationships between nodes but does consider some information from the past to regularise state variables (Equation 1; system equation). 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).

Even though the main focus of this paper is on data from fMRI, 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).

Conclusions

Our results demonstrate that DGM can 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. 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 provides news insights into directed dynamic relationships in functional connectivity.

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.
- Barnett, L., Seth, A.K., 2017. Detectability of Granger causality for subsampled continuous-time neurophysiological processes. *J. Neurosci. Methods* 275, 93–121.
- 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.
- Bullmore, E., Sporns, O., 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat. Rev. Neurosci.* 10, 186–198.
- 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., 2015a. Searching Multiregression Dynamic Models of Resting-State fMRI Networks Using Integer Programming. *Bayesian Anal.* 10, 441–478.
- Costa, L., Smith, J., Nichols, T., Cussens, J., Duff, E.P., Makin, T.R., 2015b. 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., Ushey, K., Kou, Q., Russel, N., Chambers, J., Bates, D., 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.
- Eichenbaum, H., 2017. Prefrontal-hippocampal interactions in episodic memory. *Nat. Rev. Neurosci.* 18, 547–558.
- Fornito, A., Zalesky, A., Breakspear, M., 2015. The connectomics of brain disorders. *Nat. Rev. Neurosci.* 16, 159–172.
- Fornito, A., Zalesky, A., Bullmore, E., 2016. *Fundamentals of Brain Network Analysis*. Academic Press.
- 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., Litvak, V., Oswal, A., Razi, A., Stephan, K.E., van Wijk, B.C.M., Ziegler, G., Zeidman, P., 2016. Bayesian model reduction and empirical Bayes for group (DCM) studies. *Neuroimage* 128, 413–431.
- 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.
- Gates, K.M., Molenaar, P.C.M., 2012. Group search algorithm recovers effective connectivity maps for individuals in homogeneous and heterogeneous samples. *Neuroimage* 63, 310–319.
- 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.
- Harrison, P.J., Stevens, C.F., 1976. Bayesian Forecasting. *J. R. Stat. Soc. Series B Stat. Methodol.* 38, 205–247.
- 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.
- Hyvärinen, A., Smith, S.M., 2013. Pairwise Likelihood Ratios for Estimation of Non-Gaussian Structural Equation Models. *J. Mach. Learn. Res.* 14, 111–152.
- Jin, J., Maren, S., 2015. Prefrontal-Hippocampal Interactions in Memory and Emotion. *Front. Syst. Neurosci.* 9, 170.
- 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., 2017. On the Stability of BOLD fMRI Correlations. *Cereb. Cortex* 27, 4719–4732.
- Lee, H.J., Weitz, A.J., Bernal-Casas, D., Duffy, B.A., Choy, M., Kravitz, A.V., Kreitzer, A.C., Lee, J.H., 2016. Activation of Direct and Indirect Pathway Medium Spiny Neurons Drives Distinct Brain-wide Responses. *Neuron* 91, 412–424.
- 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.
- Markov, N.T., Ercsey-Ravasz, M.M., Ribeiro Gomes, A.R., Lamy, C., Magrou, L., Vezoli, J., Misery, P., Falchier, A., Quilodran, R., Gariel, M.A., Sallet, J., Gamanut, R., Huissoud, C., Clavagnier, S., Giroud, P., Sappey-Marini, D., Barone, P., Dehay, C., Toroczkai, Z., Knoblauch, K., Van Essen, D.C., Kennedy, H., 2014. A weighted and directed interareal connectivity matrix for macaque cerebral cortex. *Cereb. Cortex* 24, 17–36.
- Menon, V., 2011. Large-scale brain networks and psychopathology: a unifying triple network model. *Trends Cogn. Sci.* 15, 483–506.
- 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., 2009a. *Dynamic Linear Models with R*. Springer New York.
- Petris, G., Petrone, S., Campagnoli, P., 2009b. Dynamic linear models, in: *Dynamic Linear Models with R, Use R*. Springer, New York, NY, pp. 31–84.
- Preti, M.G., Bolton, T.A., Van De Ville, D., 2017. The dynamic functional connectome: State-of-the-art and perspectives. *Neuroimage* 160, 41–54.
- 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., Glymour, C., 2011. Multi-subject search correctly identifies causal connections and most causal directions in the DCM models of the Smith et al. simulation study. *Neuroimage* 58, 838–848.
- 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.
- Ramsey, J.D., Sanchez-Romero, R., Glymour, C., 2014. Non-Gaussian methods and high-pass filters in the estimation of effective connections. *Neuroimage* 84, 986–1006.
- Razi, A., Kahan, J., Rees, G., Friston, K.J., 2015. Construct validation of a DCM for resting state fMRI. *Neuroimage* 106, 1–14.
- Razi, A., Seghier, M.L., Zhou, Y., McColgan, P., Zeidman, P., Park, H.-J., Sporns, O., Rees, G., Friston, K.J., 2017. Large-scale DCMs for resting-state fMRI. *Network Neuroscience* 1, 222–241.
- Ryali, S., Supekar, K., Chen, T., Menon, V., 2011. Multivariate dynamical systems models for estimating causal interactions in fMRI. *Neuroimage* 54, 807–823.
- Seghier, M.L., Friston, K.J., 2013. Network discovery with large DCMs. *Neuroimage* 68, 181–191.
- Sestieri, C., Corbetta, M., Romani, G.L., Shulman, G.L., 2011. Episodic memory retrieval, parietal cortex, and the default mode network: functional and topographic analyses. *J. Neurosci.* 31, 4407–4420.
- 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. <https://doi.org/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., Ugurbil, 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., 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., Ugurbil, 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.
- Stephan, K.E., Weiskopf, N., Drysdale, P.M., Robinson, P.A., Friston, K.J., 2007. Comparing hemodynamic models with DCM. *Neuroimage* 38, 387–401.
- 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., Abeyesuriya, 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*. <https://doi.org/10.1016/j.neuroimage.2017.06.077>
- Weilbacher, R.A., Gluth, S., 2016. The Interplay of Hippocampus and Ventromedial Prefrontal Cortex in

- Memory-Based Decision Making. *Brain Sci* 7. <https://doi.org/10.3390/brainsci7010004>
- 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

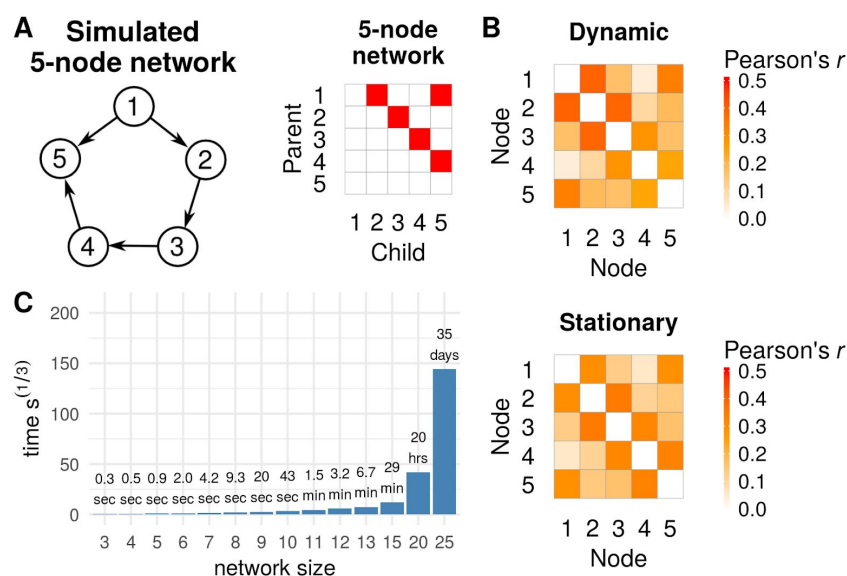


Figure 1. (A) True 5-node network and corresponding adjacency matrix. **(B)** Mean Pearson correlations across 50 simulations for the dynamic (Sim22, top) and stationary network simulations (Sim1, bottom). **(C)** Computational time estimates for a single node of a given network size for a time series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

Figure 2

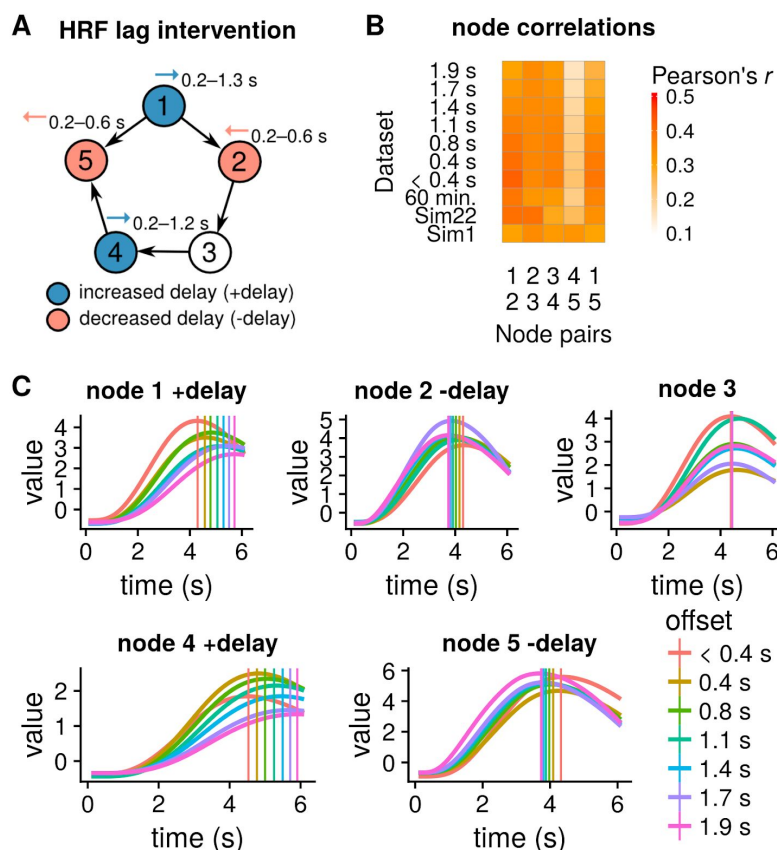


Figure 2. (A) Systematic interventions at four of the five nodes with either increased (node 1 and 4) or decreased HRF delay (node 2 and 5). We created seven simulations with a total offset between two nodes from 0.4–1.9 s **(B)** Pairwise node correlations in the different simulations. **(C)** Signal output at the 5 nodes showing the BOLD amplitude for the seven interventions where different DCM forward model parameters were used to shift peaks of the HR response, with colors from red (< 0.4 s offset) to pink (1.9 s; maximal intervention). Vertical lines indicate average time to peak after stimulus onset for each intervention type.

Figure 3

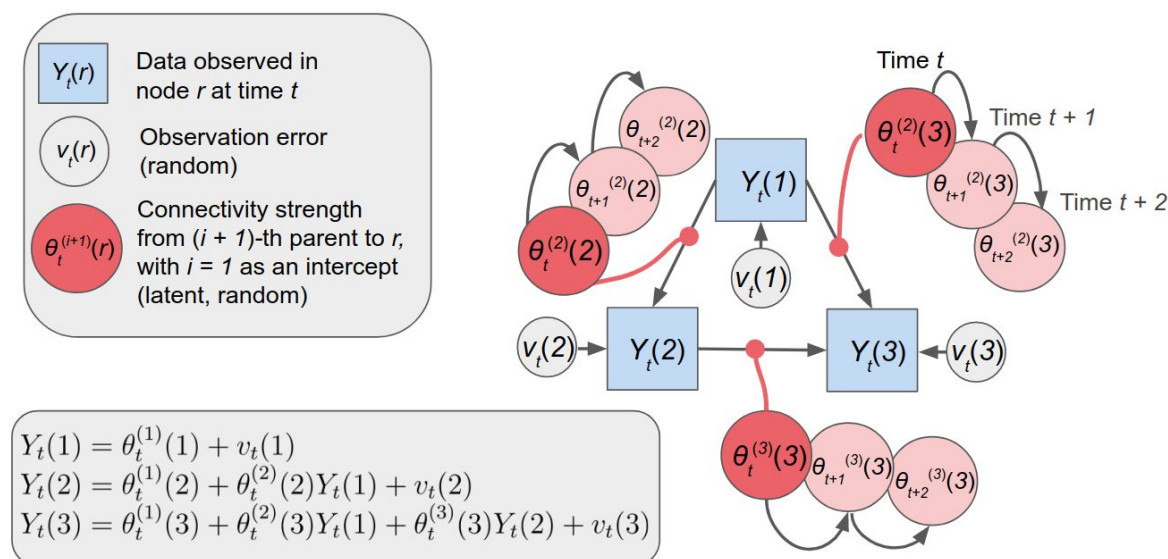


Figure 3. Dynamic Graphical Model (DGM) is a graphical model with directed relationships. $\theta_t^{(i)}(r)$ are the time-varying connectivity strength of parent region $i+1$ on child node r ($i=1$ is the intercept), or in other words, the varying regression coefficients from a dynamic regression. In this 3-node network, the model equations for the three nodes are written in the lower left box, showing the DGM to be a collection of dynamic linear models (DLM).

Figure 4




Figure 4. Pruning method's model comparison approach, illustrated on a 3 node network. Each DGM with a bidirectional edge (red; here between node A and B) is a model-pair composed of two DLMs, one for node A (parents of A) and one for B (parents of B). We compare the model evidence (ME) between three models, the bidirectional (red) m₁ and two unidirectional models (m₂ and m₃; green) where one or the other of the edges has been removed. We could choose a simpler model, and select the better unidirectional model if the difference in model evidence does not exceed an a priori defined value (in terms of a log Bayes factor).

Figure 5




Figure 5. Sensitivity, specificity, and accuracy of the methods DGM, Patel and Lingam for dynamic **(A)** and stationary simulations **(B)**. **(C)** Proportions (top) and significant proportions (below, binomial test, 5% FDR threshold) for dynamic network simulations. **(D)** same as **(C)** for stationary network simulations.

Figure 6




Figure 6. Sensitivity, specificity and accuracy for DGM ($e = 20$) for different node offsets (from 0.4 s up to 1.9 s total offset between the hemodynamic response of two nodes).

Figure 7

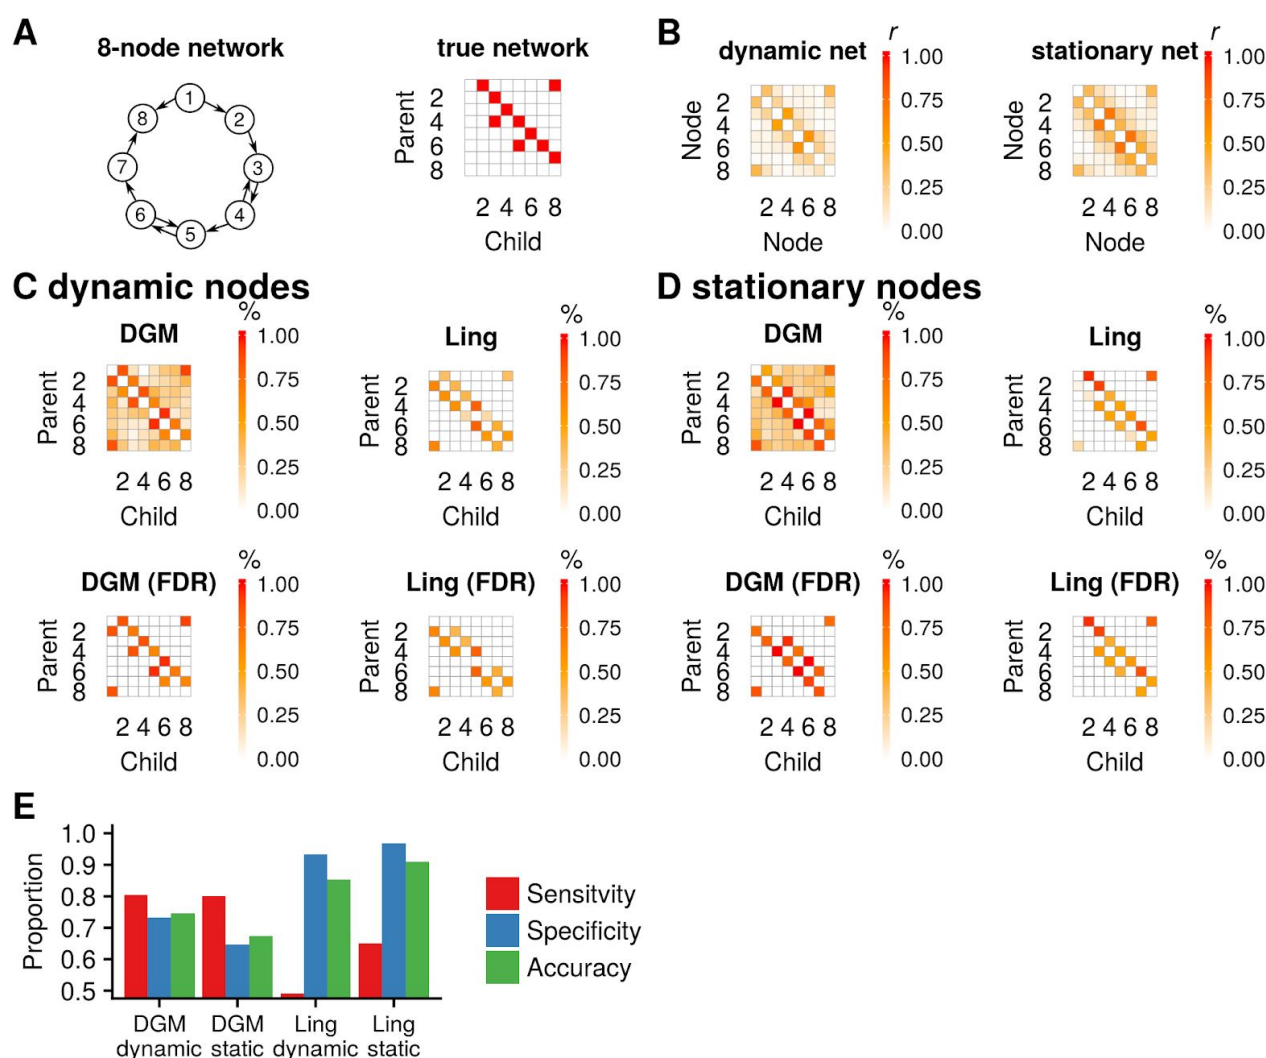


Figure 5. (A) The 8-node network with two reciprocal connections. **(B)** Pair-wise correlation matrix for both networks. **(C)** Proportions (top) and significant proportions (below, binomial test, 5% FDR threshold) for DGM and Lingam with dynamic network simulations. **(D)** same as **(C)** for stationary network simulations. **(E)** Sensitivity, specificity, and accuracy of the two network estimation methods for both types of simulations.

Figure 8

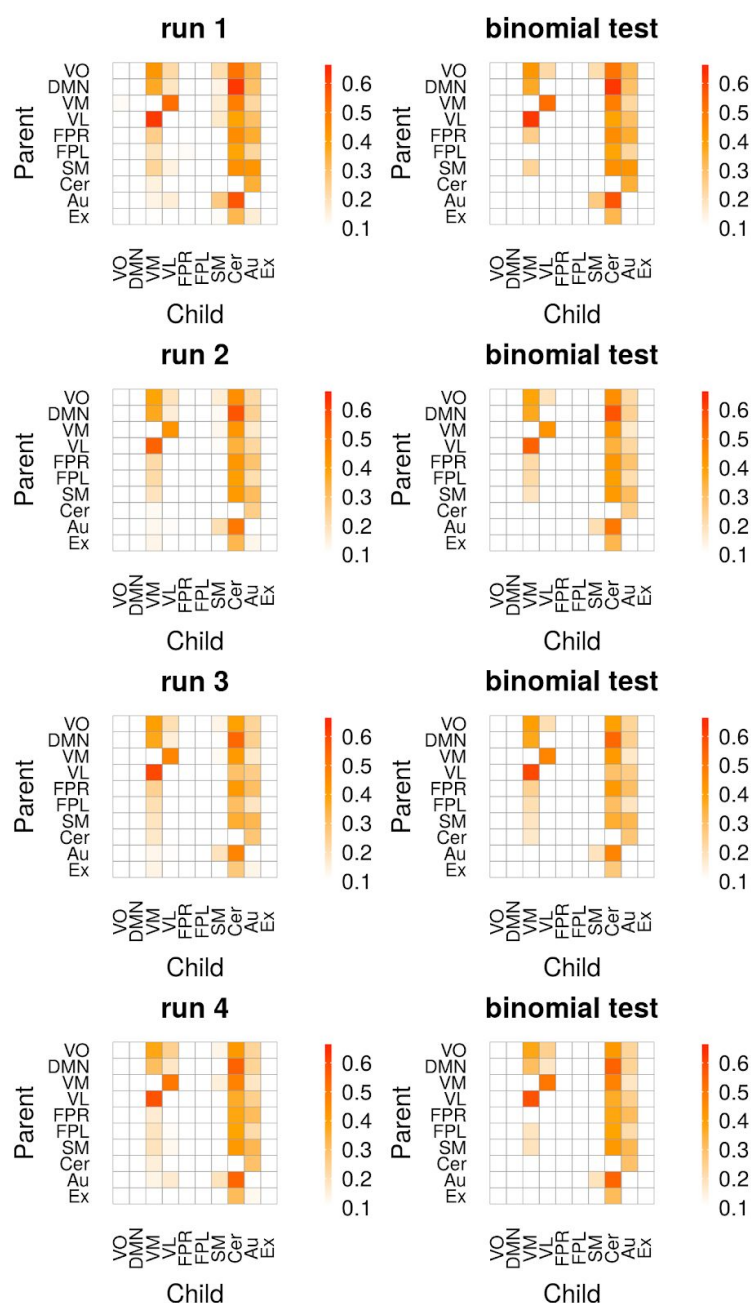


Figure 8. Proportions of edges in participants ($N = 500$) from 4 runs (top to bottom) during resting-state of approximately 15 minutes. Left shows all proportions, right shows proportions for significant edges (binomial test, 5% FDR threshold). Legend: VO, visual occipital pole; DMN, default mode; VM, visual medial; VL, visual lateral/ventral; FPR, frontoparietal right; FPL, frontoparietal left; SM, sensorimotor; Cer, cerebellum; Au, auditory; Ex, executive control (including thalamus).

Figure 9

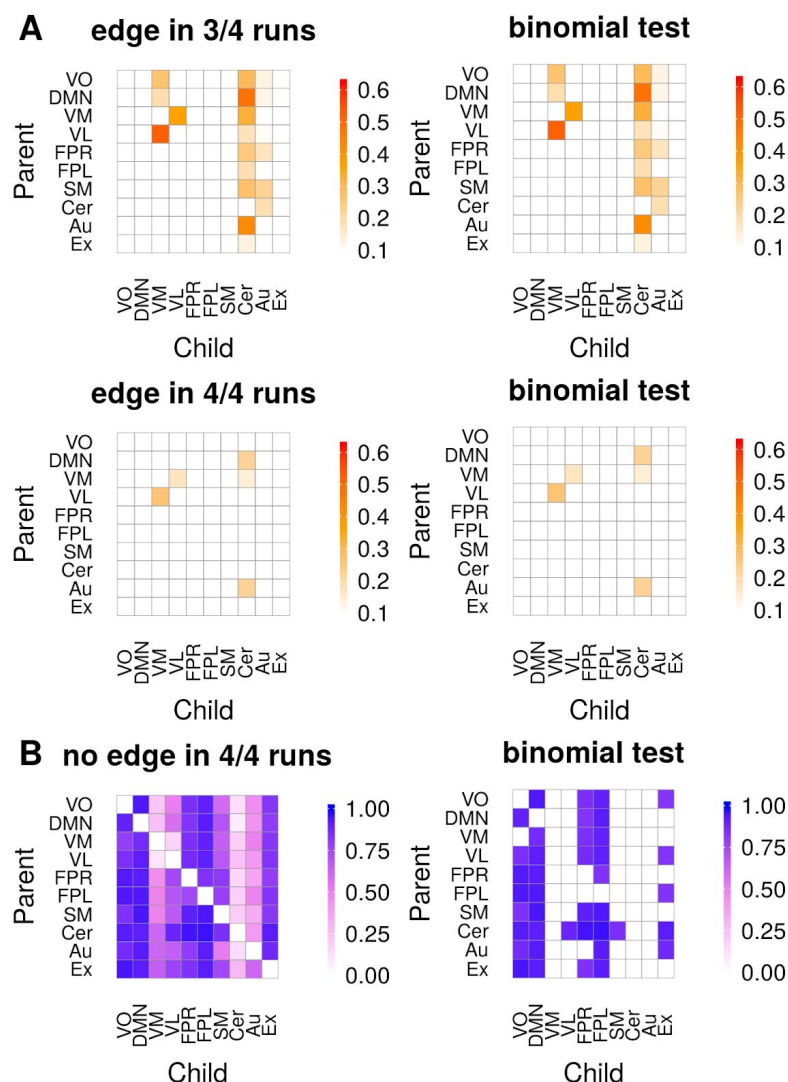


Figure 9. (A) Proportions of edges (red) from 500 participants that reoccur in three of the four runs, or in all the four runs (within-subject reproducibility). **(B)** Proportion of participants that had no edge (blue) across all runs. Left shows all proportions, right shows proportions only for significant edges (binomial test, 5% FDR threshold).

Figure 10




Figure 10. Example of time-varying connectivity estimates $\theta_i(r)$ showing the influence of the default mode network (A) and the auditory network (B) on the cerebellar network.

Figure 11

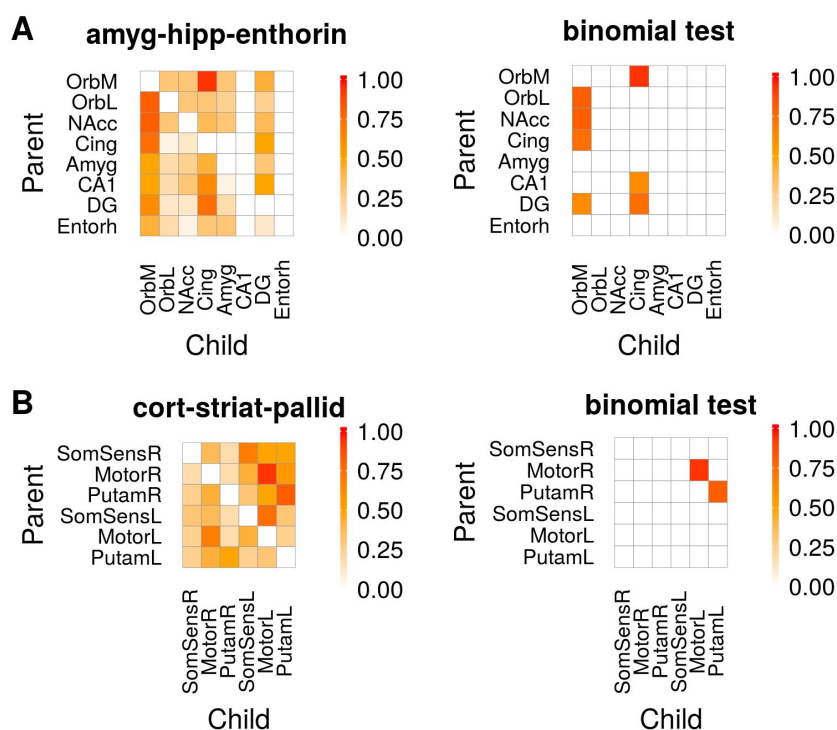


Figure 11. Proportion of edges in two functional networks of mice ($N = 16$). **(A)** amygdala-hippocampal-entorhinal network and **(B)** an cortico-striatal network. Left shows all proportions, right shows significant proportions of edges (binomial test, 5% FDR threshold). Legend: orbital area, medial (OrbM) and lateral part (OrbL); nucleus accumbens (NAcc); anterior cingulate area, dorsal part (Cing); lateral amygdalar nucleus (Amyg); cornu ammonis area 1 (CA1); dentate gyrus (DG); entorhinal area, lateral part (Entorh); primary somatosensory area, barrel field right and left (SomSensR, SomSensL); primary motor area right and left (MotorR, MotorL); caudoputamen right and left (PutamR, PutamL).

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	Number of nodes	Neural lag	Dynamic network strengths	Duration	TR	Noise	HRF lag variability (s)
Sim1	5	50 ms	No	10 min.	3 s	1.0%	$\pm 0.5^1$
Sim22	5	50 ms	Yes	10 min.	3 s	0.1%	$\pm 0.5^1$
Offset < 0.4 s	5	50 ms	Yes	10 min.	2 s	0.1%	$\pm 0.3^1$
Offset 0.4 s	5	50 ms	Yes	10 min.	2 s	0.1%	+0.21/-0.20 ²
Offset 0.8 s	5	50 ms	Yes	10 min.	2 s	0.1%	+0.48/-0.35 ²
Offset 1.1 s	5	50 ms	Yes	10 min.	2 s	0.1%	+0.69/-0.46 ²
Offset 1.4 s	5	50 ms	Yes	10 min.	2 s	0.1%	+0.85/-0.54 ²
Offset 1.7 s	5	50 ms	Yes	10 min.	2 s	0.1%	+1.04/-0.58 ²
Offset 1.9 s	5	50 ms	Yes	10 min.	2 s	0.1%	+1.24/-0.59 ²
60 min.	5	50 ms	Yes	60 min.	2 s	0.1%	$\pm 0.3^1$
8-nodes ST	8	500 ms	No	10 min.	2 s	0.1%	$\pm 0.5^1$
8-nodes NS	8	500 ms	Yes	10 min.	2 s	0.1%	$\pm 0.5^1$

¹ standard deviation of HRF peak.

² average increase (node 1 and 4) / average decrease (node 4 and 5) of the hemodynamic delay.