

RESEARCH ARTICLE

From correlation to causation: Estimating effective connectivity from zero-lag covariances of brain signals

Jonathan Schiefer^{1,2}, Alexander Niederbühl^{1,2}, Volker Pernice^{1,2}, Carolin Lennartz^{3,4}, Jürgen Hennig^{3,4}, Pierre LeVan^{3,4}, Stefan Rotter^{1,2*}

1 Bernstein Center Freiburg, Freiburg, Germany, **2** Faculty of Biology, University of Freiburg, Freiburg, Germany, **3** Department of Radiology, Medical Physics, Medical Center - University of Freiburg, Freiburg, Germany, **4** Faculty of Medicine, University of Freiburg, Freiburg, Germany

* stefan.rotter@biologie.uni-freiburg.de



OPEN ACCESS

Citation: Schiefer J, Niederbühl A, Pernice V, Lennartz C, Hennig J, LeVan P, et al. (2018) From correlation to causation: Estimating effective connectivity from zero-lag covariances of brain signals. *PLoS Comput Biol* 14(3): e1006056. <https://doi.org/10.1371/journal.pcbi.1006056>

Editor: Saad Jbabdi, Oxford University, UNITED KINGDOM

Received: September 12, 2017

Accepted: February 26, 2018

Published: March 26, 2018

Copyright: © 2018 Schiefer et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All data analysis and simulation software files are available at <https://github.com/jonathanschiefer/connFinder>. The fast fMRI data files are available from the Open Science Framework at <https://osf.io/52mf4/>.

Funding: Supported by the DFG (grant EXC 1086). The HPC facilities are funded by the state of Baden-Württemberg through bwHPC and DFG grant INST 39/963-1 FUGG. The article processing charge was covered by the open access publication fund of the University of Freiburg. The funders had no role in

Abstract

Knowing brain connectivity is of great importance both in basic research and for clinical applications. We are proposing a method to infer directed connectivity from zero-lag covariances of neuronal activity recorded at multiple sites. This allows us to identify causal relations that are reflected in neuronal population activity. To derive our strategy, we assume a generic linear model of interacting continuous variables, the components of which represent the activity of local neuronal populations. The suggested method for inferring connectivity from recorded signals exploits the fact that the covariance matrix derived from the observed activity contains information about the existence, the direction and the sign of connections. Assuming a sparsely coupled network, we disambiguate the underlying causal structure via L^1 -minimization, which is known to prefer sparse solutions. In general, this method is suited to infer effective connectivity from resting state data of various types. We show that our method is applicable over a broad range of structural parameters regarding network size and connection probability of the network. We also explored parameters affecting its activity dynamics, like the eigenvalue spectrum. Also, based on the simulation of suitable Ornstein-Uhlenbeck processes to model BOLD dynamics, we show that with our method it is possible to estimate directed connectivity from zero-lag covariances derived from such signals. In this study, we consider measurement noise and unobserved nodes as additional confounding factors. Furthermore, we investigate the amount of data required for a reliable estimate. Additionally, we apply the proposed method on full-brain resting-state fast fMRI datasets. The resulting network exhibits a tendency for close-by areas being connected as well as inter-hemispheric connections between corresponding areas. In addition, we found that a surprisingly large fraction of more than one third of all identified connections were of inhibitory nature.

study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

Author summary

Changes in brain connectivity are considered an important biomarker for certain brain diseases. This directly raises the question of accessibility of connectivity from measured brain signals. Here we show how directed effective connectivity can be inferred from continuous brain signals, like fMRI. The main idea is to extract the connectivity from the inverse zero-lag covariance matrix of the measured signals. This is done using L^1 -minimization via gradient descent algorithm on the manifold of unitary matrices. This ensures that the resulting network always fits the same covariance structure as the measured data, assuming a canonical linear model. Applying the estimation method on noise-free covariance matrices shows that the method works nicely on sparsely coupled networks with more than 40 nodes, provided network interaction is strong enough. Applying the estimation on simulated Ornstein-Uhlenbeck processes supposed to model BOLD signals demonstrates robustness against observation noise and unobserved nodes. In general, the proposed method can be applied to time-resolved covariance matrices in the frequency domain (cross-spectral densities), leading to frequency-resolved networks. We are able to demonstrate that our method leads to reliable results, if the sampled signals are long enough.

Introduction

The networks of the brain are key to understanding its function and dysfunction [1]. Depending on the methods employed to assess structure and to record activity, networks may be defined at different levels of resolution. Their nodes may be individual neurons, linked by chemical or electrical synapses. Alternatively, nodes may also be conceived as populations of neurons, with links represented by the net effect of all synaptic connections that exist between two populations. In any case, this defines the structural substrate of brain connectivity, representing the physical (causal) basis of neuronal interactions. Nodes in a brain network influence each other by sending signals. For example, the activities of nodes in a network are generally not independent, and neuronal dynamics are characterized by correlations among the nodes involved in the network. This suggests an alternative perspective on active brain networks: Functional connectivity assigns a link to a pair of nodes to the degree to which their activities are correlated. It has been argued that this concept emphasizes connections that “matter”, including the possibility that the same substrate may give rise to different networks, depending on how they are used. As a consequence, functional connectivity and structural connectivity are not equivalent. A well-known phenomenon is that two nodes may be correlated, even if there is no direct anatomical link between them. For example, a shared source of input to both nodes may generate such a correlation, which does not correspond to a direct interaction between the two nodes. Apart from that, correlation is a symmetric relation between two nodes, whereas a physical connection implies a cause-effect relation that is directed. There have, in fact, been multiple attempts to overcome the shortcomings of functional connectivity, especially the lack of directed interaction. The term *effective connectivity* has been suggested for this [2]. The idea is to bring the networks, inferred from activity measurements, closer to structural connectivity, which can only be inferred with anatomical methods. The dichotomy between structural and functional aspects of connectivity raises the general question whether it is possible to infer brain networks from recorded activity. We are only beginning to understand the forward link between structural connectivity and functional connectivity. As a consequence, it is possible to compute correlations from connectivity in certain simplified network

scenarios [3]. The correspondence between connectivity and correlation, however, is not one-to-one. Networks with different connectivity may lead to exactly the same correlations between nodes. As a consequence, the inverse problem of inferring connectivity from correlation is generally ill-defined. As we will demonstrate in this paper, additional assumptions about the connectivity can help to resolve the ambiguity. Specifically, we search for the network with the lowest number of nonzero edges (via L^1 -minimization) to disambiguate the problem. Structural, functional and effective connectivity are not equally well accessible. Some aspects of the anatomical structure can be assessed *post mortem* by invasive tracing methods, or non-invasively by Diffusion Tensor Imaging, DTI. In contrast, functional connectivity is based on statistical relationships between the activity of neuronal populations and can be easily estimated from recorded signals. For estimating effective connectivity there are methods like Dynamic Causal Modelling, DCM [4, 5], Granger causality [6] and others [7–13]. Only few methods to infer effective connectivity, however, can deal with large numbers of nodes (40 or more) based on zero-lag correlation only. However, they are either limited to small networks [14], or to directed acyclic graphs [15]. Here, we are proposing a new method for the estimation of effective connectivity from population activity in the brain, especially BOLD-related signals. The new method is a variant of the procedure described in [16], based on a L^1 -minimization. For the method proposed here it is sufficient to use zero-lag covariances to estimate directed effective connectivity.

Materials and methods

Estimation method

The main idea of our estimation method is inspired by the finding, “that the key to determining the direction of the causal relationship between X and Y lies in ‘the presence of a third variable Z that correlates with Y but not with X ,’ as in the collider $X \rightarrow Y \leftarrow Z$...” [17, 18]. Similarly, assuming a linear interaction model, the presence of a collider structure in a network (see Fig 1) produces specific entries in the corresponding inverse covariance (precision) matrix. Fig 1 shows a disconnected network in the left column, and a network which induces the same covariance matrix if all links have opposite direction in the middle column. In the latter case an estimation of the direction is impossible, because there is simply no information about it in the covariance matrix. Whenever a collider structure is present, however, the entry in the inverse covariance matrix for the two source nodes (here, 2 and 3) is non-zero. This is due to the fact that in a linear model the entry in the inverse covariance matrix depends not only on the connections of the nodes 2 and 3, but also whether these nodes have a common target. This means the presence of a collider structure allows us to disambiguate the direction of this particular connection.

We consider here a scenario, where the interaction between nodes is described by a generic linear model. Assuming stationarity, let the neural activity $x(t)$ be implicitly defined by the consistency equation

$$x(t) = (G * x)(t) + v(t) \tag{1}$$

where $G(t)$ is a matrix of causal interaction kernels and $v(t)$ denotes fluctuating external inputs (“driving noise”). All variables are also listed in Table 1. Fourier transformation of Eq (1) yields

$$\hat{x}(f) = \hat{G}(f)\hat{x}(f) + \hat{v}(f)$$




Fig 1. Collider structures are encoded in the inverse covariance matrix. Upper row: Three simple network architectures. Lower row: The corresponding inverse covariance matrices, red color represents positive entries, blue color stands for negative ones. In the left and middle column, the entries (2, 3) and (3, 2) are 0. The only difference between the right column and the middle column is that the connection between node 1 and 2 is flipped, such that nodes 1, 2 and 3 form a collider structure. Although there is still only an indirect connection between node 2 and 3, the entry in the corresponding inverse covariance matrix is now non-zero.

<https://doi.org/10.1371/journal.pcbi.1006056.g001>

and simple rearrangement leads to

$$\hat{x}(f) = (\mathbb{1} - \hat{G}(f))^{-1} \hat{v}(f)$$

where \hat{x} denotes the Fourier transform of x . The cross spectral density of the signals is then given by

$$\hat{C}(f) = (\mathbb{1} - \hat{G}(f))^{-1} \hat{Z}(f) (\mathbb{1} - \hat{G}^*(f))^{-1}$$

where $\hat{Z}(f)$ is the cross-spectral density of the external inputs. It follows

$$\begin{aligned} \hat{C}^{-1}(f) &= (\mathbb{1} - \hat{G}^*(f)) \hat{Z}^{-1}(f) (\mathbb{1} - \hat{G}(f)) \\ &= B^*(f) B(f) \end{aligned} \tag{2}$$

Table 1. Variables used for the estimation method and simulation.

Variable name	Symbol
node activity	$x(t)$
network connectivity	G
external inputs	$v(t)$
cross-spectral density	$\hat{C}(f)$
covariance of external input	$\hat{Z}(f)$
unitary matrix	U
L^1 -norm cost function	Γ
gradient	d
initial matrix	B_0
Wiener process	$w(t)$
stationary covariance matrix	σ
simulation step	Δt
time constant of activity	τ
regularisation-controlling parameter for regularized ICOV	λ

<https://doi.org/10.1371/journal.pcbi.1006056.t001>

with $B(f) = \sqrt{\hat{Z}^{-1}(f)} (\mathbb{1} - \hat{G}(f))$. In our model, we assume that the components of the external fluctuating input are pairwise stochastically independent for all nodes. Then, \hat{Z} is a diagonal matrix, and we make the additional assumption that $\hat{Z} = \mathbb{1}$. For the linear model considered here, there is a relation between covariance and connectivity, which can be exploited for the estimation of connectivity from correlation. In the case $\hat{Z} = \mathbb{1}$ it is given by

$$\hat{C}^{-1}(f) = (\mathbb{1} - \hat{G}^*(f)) (\mathbb{1} - \hat{G}(f)) = \mathbb{1} - \hat{G}(f) - \hat{G}^*(f) + \hat{G}^*(f)\hat{G}(f)$$

where the last term contributes the information of the collider structures. If the matrix product $\hat{G}^*(f)\hat{G}(f)$ has a non-zero off-diagonal entry the corresponding nodes have outgoing connections terminating at the same node, which means these nodes form a collider. It is clear that for any unitary matrix $U \in \mathcal{U}(n)$ the product UB is still a solution of Eq (2), as $U^*U = \mathbb{1}$. We will resolve this ambiguity with an L^1 minimization which is known to prefer sparse solutions under certain conditions [19]. In order to find G from a given C we first fix an initial matrix B , and then search for a unitary matrix $U \in \mathcal{U}(n)$ such that $\|UB\|_1$ is minimal, so we are minimizing the function

$$\begin{aligned} \Gamma : \mathcal{U}(n) &\longrightarrow \mathbb{R} \\ U &\longmapsto \|UB\|_1. \end{aligned} \tag{3}$$

Gradient descent. To estimate the connectivity matrix from the covariance matrix we use a conjugate gradient descent algorithm similar to [20, 21] for minimizing the function $\Gamma(U)$ given in Eq (3), implemented in Python. For details please see S1 Text. For the gradient

$$d_{ij} = \frac{\partial \Gamma(U)}{\partial U_{ij}}$$

of the cost function $\Gamma(U)$, $a = \frac{1}{2}(d - d^*)$ is skew-hermitian, and the matrix exponential of a skew-hermitian matrix is unitary. This means, starting in a point U_{act} and choosing an appropriate step size δ , we obtain a point $U_{new} = \exp(-\delta a)U_{act}$ with $\Gamma(U_{new}) < \Gamma(U_{act})$. In other words, the new point has a smaller L^1 -norm than the old one and still satisfies the condition $\hat{C}^{-1} = B^*U_{new}^*U_{new}B$. Iterating this procedure until convergence leads to a point with locally minimal L^1 -norm. The two conditions for convergence are inspired by [21]. The first one is a condition on the norm of the gradient. In each step, it is checked if

$$\|d - Ud^*U\|_F = \sqrt{\sum_{ij} |d_{ij}(Ud^*U)_{ij}|^2} < \text{gtol}$$

is fulfilled, where $\|\dots\|_F$ is the Frobenius norm and $\text{gtol} > 0$ is the convergence tolerance. As a second (alternative) condition, it is checked whether simultaneously

$$\frac{\|U - U_{old}\|_F}{\sqrt{N}} < \text{xtol} \quad \text{and} \quad \frac{|\Gamma(U_{old}) - \Gamma(U)|}{|\Gamma(U_{old} + 1)|} < \text{ftol}$$

are fulfilled. The values used are listed in Table 2. Before convergence the cost function typically oscillates around a certain value. To avoid stopping at a random phase of this oscillation, as a final step we apply a line-search, for details see S1 Text. The described gradient descent algorithm provides an efficient way for minimizing Eq (3). When calculating the gradient, we neglect the diagonal. Consequently, we also neglect the diagonal of the resulting estimated matrix, so we are not able to study self connections of the nodes.

Table 2. Parameter used for estimation.

Parameter	Value
xtol	$0.7 \cdot 10^{-2}$
ftol	$0.7 \cdot 10^{-4}$
gtol	$0.7 \cdot 10^{-2}$
κ	500
λ	5

<https://doi.org/10.1371/journal.pcbi.1006056.t002>

Initial condition. As starting condition for the gradient descent we use a matrix $B_0(f)$ such that

$$B_0^*(f)B_0(f) = \hat{C}^{-1}(f).$$

There are many ways to choose a B_0 with this property, we found the following choice efficient: As \hat{C} is the cross-spectral density it is positive definite, and so is \hat{C}^{-1} . Thus, there is exactly one positive definite square root of \hat{C}^{-1} [22] which can be calculated by

$$B_0 = \sqrt{\hat{C}^{-1}} = W\sqrt{E}W^* \tag{4}$$

where the columns of W are the eigenvectors of \hat{C}^{-1} , and E is the matrix with the corresponding eigenvalues of \hat{C}^{-1} on the diagonal. Thus we initialize the gradient descent with $U_0 = \mathbf{1}$ and B_0 given by Eq (4).

Step size selection. A critical part of the optimization is the selection of an appropriate step size. If the step size is too large, one might miss the minimum. If the step size is too small, the optimization converges very slowly. For the gradient descent, we use an adaptive scheme inspired by [20], where the step-size depends on the largest eigenvalue of the actual gradient: Let λ_{\max} be the largest eigenvalue of d_{act} , the step size is given by

$$\delta_{\text{act}} = \frac{2\pi}{|\lambda_{\max}| \cdot \kappa}$$

where κ is constant. The intuition behind that, is that smooth cost functions along a geodesic on the unitary manifold are almost periodic. So the step size should be a fraction of the period of this function. This is achieved by the scaling with the largest eigenvalue, which allows us to take a scale-invariant fraction of this period.

Validation methods

Noise-free covariance matrices. We assume that the interactions among neuronal populations can be described by a linear model, see Eq (2) with $Z = \mathbf{1}$. This model allows us to derive a relation between the connectivity matrix of the network G and the inverse cross-spectral density matrix \hat{C}^{-1} of the measured activity

$$\hat{C}^{-1} = (\mathbf{1} - \hat{G}^*)(\mathbf{1} - \hat{G}) = \mathbf{1} - \hat{G}^* - \hat{G} + \hat{G}^*\hat{G}. \tag{5}$$

Given a sampled connectivity matrix G we can calculate the inverse covariance matrix directly using Eq (5). For all simulations, half of the connections were negative (inhibitory) connections, the absolute strength was the same for all connections and 20 repetitions were simulated. As connectivity profiles we used random Erdős-Rényi networks.

Ornstein-Uhlenbeck processes. To validate our inference procedure before applying it to the network inference from measurements of neuronal activity we simulated stationary signals. Since there is no gold-standard for simulations of fMRI data [23], we based our simulations on the Ornstein-Uhlenbeck process [24], which provides a simple linear model for neural activity.

$$dx(t) = Ax(t)dt + dW(t) \tag{6}$$

where A is a matrix and W a Wiener process. In our applications, we parametrize this matrix as $A = \frac{1}{\tau}(G - \mathbb{1})$ with real-valued connectivity matrix G and time constant τ . For this process, it is possible to calculate the stationary covariance matrix Σ from the continuous Lyapunov equation

$$\mathbb{1} = A\Sigma + \Sigma A^T.$$

In fact, we simulated the process in discrete time. In analogy with [25] we use a multivariate version of the exact update formula

$$x(t + \Delta t) = e^{A\Delta t}x(t) + n(t), \tag{7}$$

where $n(t) \sim N(0, \Sigma)$ is normally distributed, with Σ being the stationary covariance matrix described above. As a final step, we filter the time series $x(t)$ with the canonical hemodynamic response function (HRF) [26, 27]. To match the data obtained in brain scans sampled at a temporal resolution of $\Delta t = 0.1$ s, we used random connectivity profiles G with a connection probability $p = 0.1$ (Erdős-Rényi model), 50% negative entries, and a spectral radius of $\rho = 0.3$. All parameters used are listed, once more, in Table 3. Before calculating the covariance C , the data is standardized such that the mean is 0 and the variance is 1 for all components of the time series. We add normally distributed observation noise u_{obs} with a $\mathcal{N}(0, \sigma_{\text{obs}})$ distribution to the simulated signal. After the simulation we calculated the signal-to-noise ratio according to

$$\text{SNR} = \frac{\sigma_x^2}{\sigma_{\text{obs}}^2}$$

where σ_x^2 denotes the variance of the signal.

Performance measures

When estimating connectivity from simulations with known underlying network structure (ground truth), one can quantify the performance of the estimation. For measuring the accuracy of our estimation we employ three different methods. First, we use the area under the ROC-curve (AUC). The ROC (receiver operating characteristic) curve is obtained as following:

Table 3. Parameter used for simulations.

Parameter	Value
repetitions	20
network type	Erdős-Rényi
N	100
p	0.1
T	350000 s
dt	0.1 s
τ	0.1 s
ρ	0.3

<https://doi.org/10.1371/journal.pcbi.1006056.t003>

For each possible parameter value (in our case the threshold for the existence of a connection), the number of true-positives (TP) and false-positives (FP) is used to calculate the true-positive rate (or recall) $TP/(TP+ FN)$ and the false-positive rate $FP/(FP+ TN)$. The ROC curve is then obtained by plotting the true-positive-rate against the false-positive rate. Secondly, we use the average precision score (PRS) which is the area under the precision-recall curve. This also includes the false-negatives (FN) (precision: $\frac{TP}{TP+FP}$). If both AUC and PRS are equal to 1, the connections in the network are perfectly estimated. Sample curves are shown in Fig 2D. Thirdly, we calculate the Pearson Correlation Coefficient (PCC) which in contrast to the measures defined before also take the strength and the sign of the interactions into account. This also means that this measure is less suited to assess whether a connection exists or not. It rather measures whether the estimated connections have the same strength as the original ones. We consider all three performance measures simultaneously to establish the quality of our estimates.

Experimental fMRI data

Seven healthy subjects underwent a 20-minute resting-state fMRI experiment on a 3 T Siemens Prisma scanner. The data was acquired using the MREG sequence [28], yielding a high temporal resolution (TR = 0.1 s, 12000 time points) that facilitates functional connectivity analyses [29]. The other sequence parameters were TE = 36 ms, FA = 25°, 64 × 64 × 50 matrix and 3mm isotropic voxel size. Additionally, cardiac and respiratory signals were recorded with the ECG and abdominal breathing band from the scanner’s physiological monitoring unit. Motion correction was done with FSL and physiological noise correction was performed with RETROICOR [30]. Average CSF and white matter signals were regressed out, but no global signal regression was performed. Following image normalization to MNI space, voxels were parcellated according to the AAL atlas (excluding the cerebellum), and the mean activity within each atlas region was calculated. The connectivity was then estimated using zero-lag covariances of the standardized signals.

Ethics statement. The experiments have been approved by the Ethics Committee of the University Medical Center Freiburg.

Results

Noise-free covariance matrices

Intrinsic properties of our new estimation procedure can be identified by studying the performance of the method for perfectly estimated (noise-free) covariance matrices. This way we address properties that do not depend on any particular feature of the underlying data, and that are not due to the success of the measurement process. In particular, we show for which types of networks our estimation procedure gives good results on technical grounds, with a wide range of networks hopefully including those arising in applications. We used random Erdős-Rényi connectivity profiles for all simulations. The macro-connectivity between neuronal populations has to satisfy certain conditions in order to be tractable by our methods. Two of these conditions concern the dynamic stability of the network and the strength of the interactions. There is a trade-off between the number of physical links and the resulting strength of macro-connections, and the dynamic stability of the network. To study the performance of our method in these various regimes, we separately varied the network size N , the connection probability p , and the absolute strength of connections $|J|$ in the connectivity matrix G , while the fraction of inhibitory couplings was kept at 50%. The spectral radius ρ of

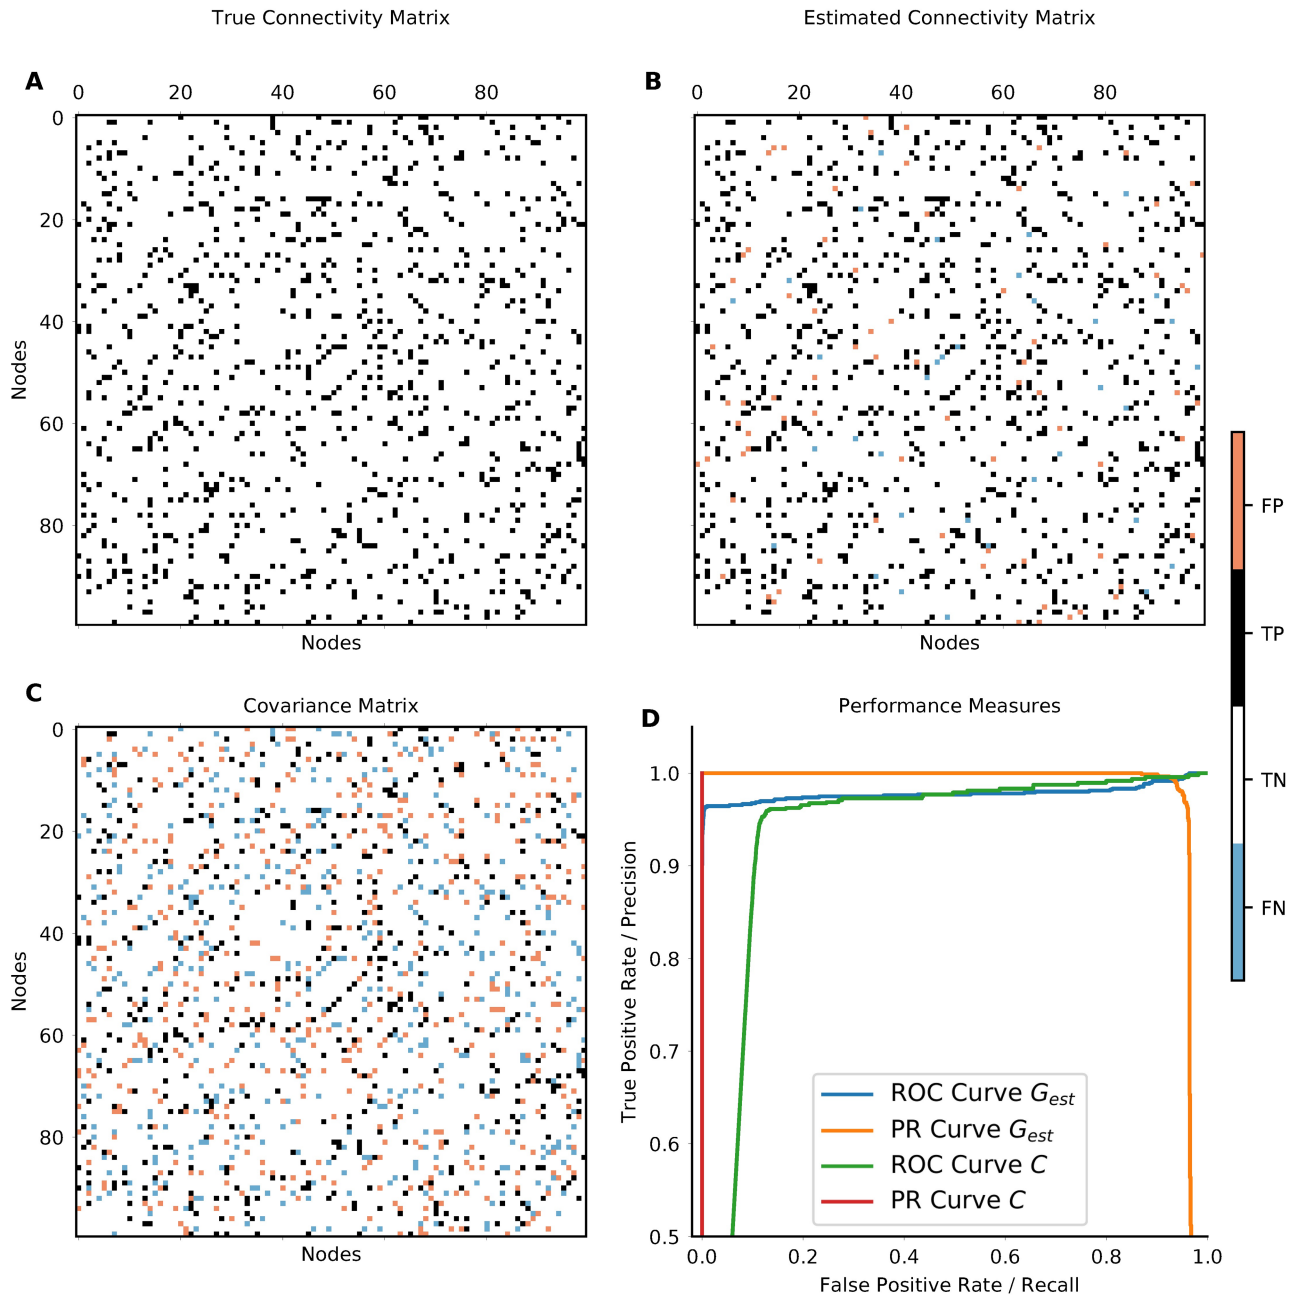


Fig 2. Networks inferred from a simulated Ornstein-Uhlenbeck process. **A** shows the original network. **B** shows the network inferred with our new method from the zero-lag covariances. White and black entries indicate true negative (TN) and true positive (TP) connections, blue and red entries indicate false negative (FN) and false positive (FP) connections, respectively. In this example, the performance measures are AUC = 0.98, PRS = 0.97 and PCC = 0.95. **C** depicts the sample covariance (functional connectivity) matrix directly estimated from the data. In **C**, as a consequence of symmetry, the number of wrongly estimated connections is quite high, the performance measures are AUC = 0.93, PRS = 0.54, and PCC = 0.29. **D** shows the Receiver Operating Characteristic Curve and the Precision Recall Curve for the networks estimated from zero-lag covariance G_{est} in blue/orange and of the functional connectivity C in red/cyan. The areas under these curves are the AUC and PRS, respectively.

<https://doi.org/10.1371/journal.pcbi.1006056.g002>

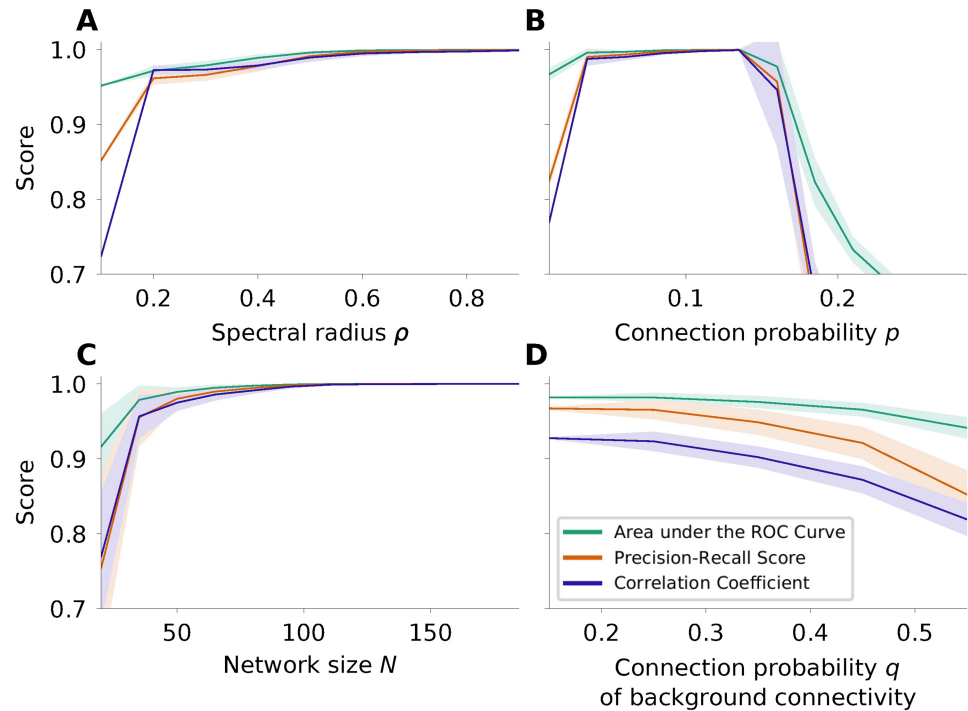


Fig 3. Effects of spectral radius ρ , connection probability p and network size N . Here we consider the case of noise-free covariance matrices, which were created based on the theory of the underlying model. The quantities considered are the area under the ROC curve (AUC; green), the precision recall score (PRS; orange) and the Pearson correlation coefficient (PCC; purple). The shaded areas indicate the mean \pm standard deviation computed over 20 realizations. **A** If the network interaction is larger than ρ_{\min} , it has relatively little effect on the performance of the estimation. Even in the extreme case, where ρ is close to 1, the estimation works well. **B** Performance of the estimation for different sparsity levels, encoded by the respective connection probabilities p . As expected, for non-sparse networks the performance of the algorithm degrades dramatically. **C** Performance of the estimation for increasing network size. Our results indicate clearly that bigger networks can be better reconstructed. Applicability may be limited by the numerical effort associated with the optimization. **D** Performance of the estimation in presence of weak background connections. It is nevertheless possible to infer the skeleton of strong connections with high fidelity.

<https://doi.org/10.1371/journal.pcbi.1006056.g003>

the bulk eigenvalue spectrum is approximately given by

$$\rho^2 = J^2 p(1 - p)N. \quad (8)$$

The default values of the parameters used in our study were $N = 100$, $p = 0.1$ and $\rho = 0.7$, where only one of them at a time was systematically varied. Low values of the spectral radius ρ correspond to networks with weak recurrent interaction and high values to networks with strong interaction, respectively. According to the model of network interaction assumed here, the networks need to have a spectral radius $\rho > 0$ for network interaction to be present and $\rho < 1$ for the dynamics to be stable. First, our results in Fig 3A indicate that a certain minimal level of interaction is necessary to be able to estimate the connections reliably. Above a value of $\rho_{\min} = 0.2$, the influence of the spectral radius on the performance of the estimation is weak, but the larger the spectral radius is the better the estimation gets.

Secondly, the connection probability of the network influences the quality of the estimation. For all connection probabilities tested here the network size was kept constant at $N = 100$ nodes. The networks were constructed such that the strength $|J|$ of all connections was the same and such that the spectral radius ρ was constant according to Eq (8). Fig 3B shows that the estimation works very well for sparse matrices with a connection probability in the range

between 5% and 15%. For networks with higher connection probability and equally strong connections, the performance decreases as expected, due to the bias associated with L^1 -minimization. But even for a connection probability of $p = 0.21$, a fraction of 14.2% of the estimated connections are false negative, and 3.3% are false positive. More than 90% of the correctly estimated connections have the correct sign. In applications, the focus of the estimation often lies on the strongest connections in the network. In networks with a background of weak connections and a sparse skeleton of stronger connections, it is possible to selectively estimate these strong links although, strictly, the assumption of a sparse network is violated. Fig 3D shows the performance of our method for such networks: the networks consist of a skeleton of strong connections with connection probability $p = 10\%$ and a connection strength derived from Eq (8) for $\rho = 0.7$. Additionally, we created a second network with weaker connections for various connection probabilities q . The two networks were combined by adding the connectivity matrices. The connection strength of this weaker connections is also derived from Eq (8), with a spectral radius of the background network being 20% of the spectral radius of the skeleton network. Then the performance of the estimation is calculated with respect to the skeleton of strong connections.

Thirdly, to be applicable to a broad range of data types, a method of connectivity estimation should perform stable for different network sizes N . For most common types of non-invasive recordings of population activity the number of nodes considered is in the range between 30 and 150. It is, of course, possible to consider larger networks, although the estimation becomes computationally more expensive. The runtime of the algorithm for networks with 200 nodes still in the range of seconds on a state-of-the-art desktop computer, but even networks with 1000 nodes or more are tractable. The strength of the connections $|J|$ are set such that the spectral radius ρ of G is constant; the connection probability is constant at $p = 0.1$. Fig 3C shows that our method performs better for bigger networks. We have observed that the L^1 cost landscape becomes smoother for larger networks.

Ornstein-Uhlenbeck processes as model for BOLD signals

In order to create surrogate data which fit fast fMRI data [28], we simulated interacting stochastic processes known as Ornstein-Uhlenbeck processes. In this case, the performance of the network inference depends on how well the inverse covariance matrix, which is the basis of the estimation, can be derived from the data. In addition to finite size effects, we studied the impact of observation noise on the performance, see Fig 4. We used $N = 100$, $p = 0.1$, $dt = 0.1$ s, $\rho = 0.74$ and $\tau = 0.1$ s as default values of the parameters. Generally, it seems natural to use Welch's method to calculate cross-spectral densities directly, and then to estimate the connectivity for each frequency band separately. For the data described here, however, we can estimate the connectivity from zero-lag sample covariances in the time domain. This is possible when the mass of the covariance function is concentrated very close around lag 0. Then lag 0 is the only one contributing to the integral of the covariance function, which corresponds to the cross-spectral density $\hat{C}(0)$. As shown in Fig 4A, with noisy data the AUC is still good, but the PRS is lower than in the case, where the covariance is known without error. However, for a signal-to-noise ratio above 1 the performance improves very quickly.

In the case of fMRI usually the whole brain is scanned, and there are no unobserved nodes in the network. However, for other data types (e.g. fNIRS) only parts of the brain can be observed. The question then is, whether this sub-network can nevertheless be reconstructed from the recorded signals. To model this scenario, we took simulated data and removed randomly a certain subset of components from the dataset. The interaction of the removed nodes is then not part of the covariance matrix of the reduced dataset, although the unobserved

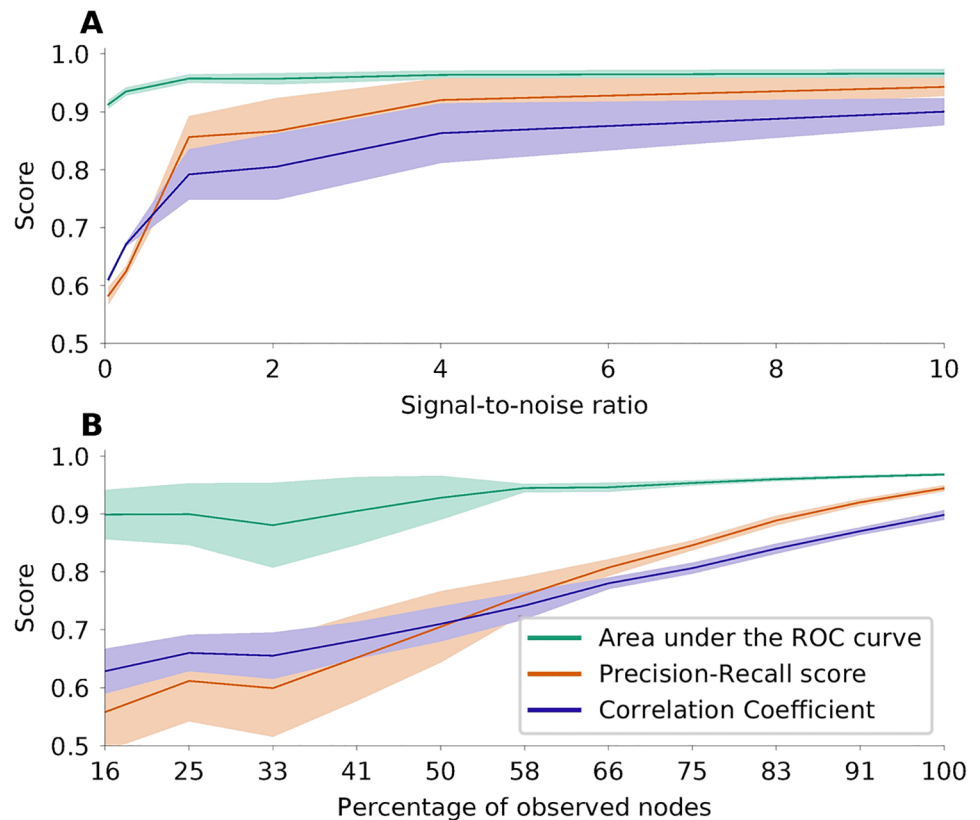


Fig 4. Performance of network inference based on simulated Ornstein-Uhlenbeck processes. Same colors as in Fig 3. **A** Performance of the estimation when measurement noise is added. **B** Performance of the estimation if only parts of the network are observable. The fraction of observed nodes in a network are indicated on the x-axis. The total number of nodes in the network was $N = 180$.

<https://doi.org/10.1371/journal.pcbi.1006056.g004>

nodes of course still exert their influence on the observed ones. The performance of the estimation of the sub-network based on the reduced dataset is shown in Fig 4B. Our analysis shows very clearly that the estimation still leads to reasonable results under these conditions. In fact, we can demonstrate that we are inferring causal connections only: For unconnected observed nodes X , Y and a latent node L connected to both X and Y , our method does not erroneously indicate a link between X and Y . One key factor for a reliable estimation of the covariance matrix is the amount of data available. This depends on the length of the measurement or simulation, and on the sampling rate. Since fast fMRI time series are obtained by measuring the BOLD response as a proxy of neuronal activity, the time scale of the measured data is relatively slow compared to the time scale of the underlying neuronal activity. Fig 5 shows the performance of network inference depending on the amount of data available, and on the time-scale of the neuronal activity. Not surprisingly, the more voluminous the dataset is, the better the estimation gets. On the other hand, it shows that the estimation generally leads to better results for slower temporal dynamics. Also, for data of sufficient length with a fairly good signal-to-noise ratio, the estimation of the connectivity is possible even when only a part of the network is observed. To allow comparison of our new method with other known methods for network inference [31–33], we applied it to the *NetSim* dataset provided by [34]. For details on the result of this, please see Fig A in S1 Text.

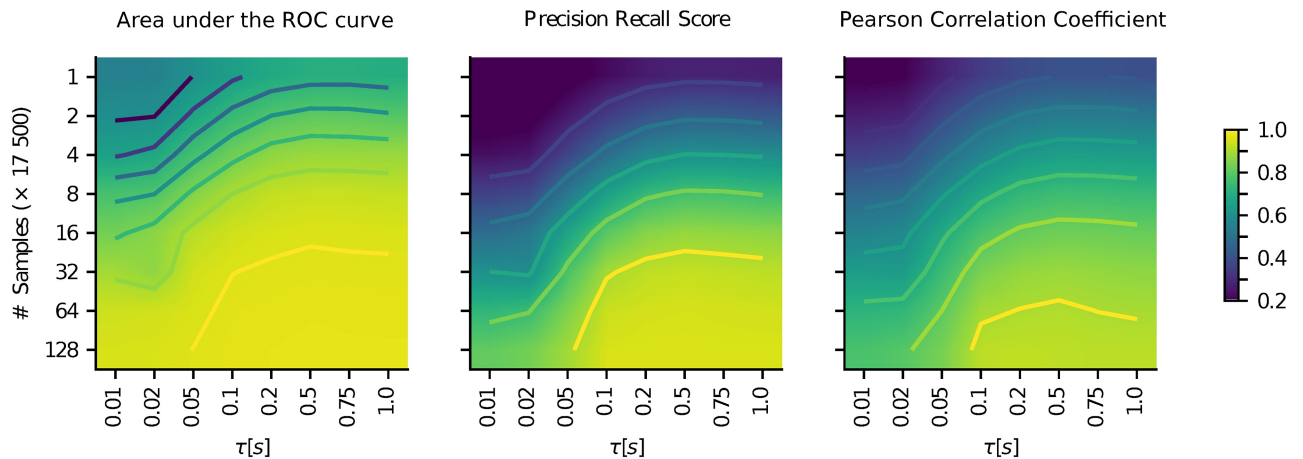


Fig 5. Performance (color coded) of the estimation depending on data length (y-axis) and time scale of the activity (x-axis). Both scales are logarithmic. For interpolation a bilinear method is used.

<https://doi.org/10.1371/journal.pcbi.1006056.g005>

fMRI data

We estimated connectivity from seven fast fMRI datasets, for details see the methods section. The resulting networks, after a threshold of 10% was applied, consist of 810 connections for each dataset. The threshold of 10% was chosen arbitrarily. In Fig B in *S1 Text* we show the histogram of estimated connection strengths for all seven reconstructed networks before thresholding. The threshold is derived from the 10% strongest connections, disregarding their signs. As there is generally no full ground truth for the connectivity inferred from human fMRI recordings available [31, 32, 35], we cannot definitely assess the degree to which the result of our inference are correct. We can, however, establish whether they are plausible. One representative connectivity matrix is shown in Fig 6. On average, 34% of the connections were inhibitory, with negligible variability across subjects. Of all connections found, 301(37%) were found in four subjects or more, and 4872 out of 8100 possible connections were absent in all subjects. On average, 245 of the connections were bi-directional and 565 connections were identified only for one direction. In general, close-by areas are more likely to be connected than more distant ones. This fact is (approximately) represented by a concentration of connections along secondary diagonals in the within-hemisphere blocks. Also, there are frequent inter-hemispheric connections between corresponding areas. This fact is represented by the diagonal entries in the across-hemisphere blocks.

Comparison with the Regularized Inverse Covariance (RIC) method

As mentioned above, different heuristics have been suggested to reconstruct networks from neuronal signals. In Fig 7 we compare the performance of the new method we propose here and the established method of Regularized Inverse Covariance [34], based on the implementation provided at <https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FSLNets>. Our comparison clearly shows that our new method performs significantly better than the Regularized Inverse Covariance method, mainly, because the latter cannot establish the direction of connections. The superior performance of the new method is reflected in higher values for all three performance measures, in particular PRS and PCC. As regularization parameter required by the software toolbox, we used $\lambda = 5$.

Furthermore, we applied the RIC method on all seven MREG datasets described before. A threshold was applied, such that only the 10% strongest connections are retained. To compare

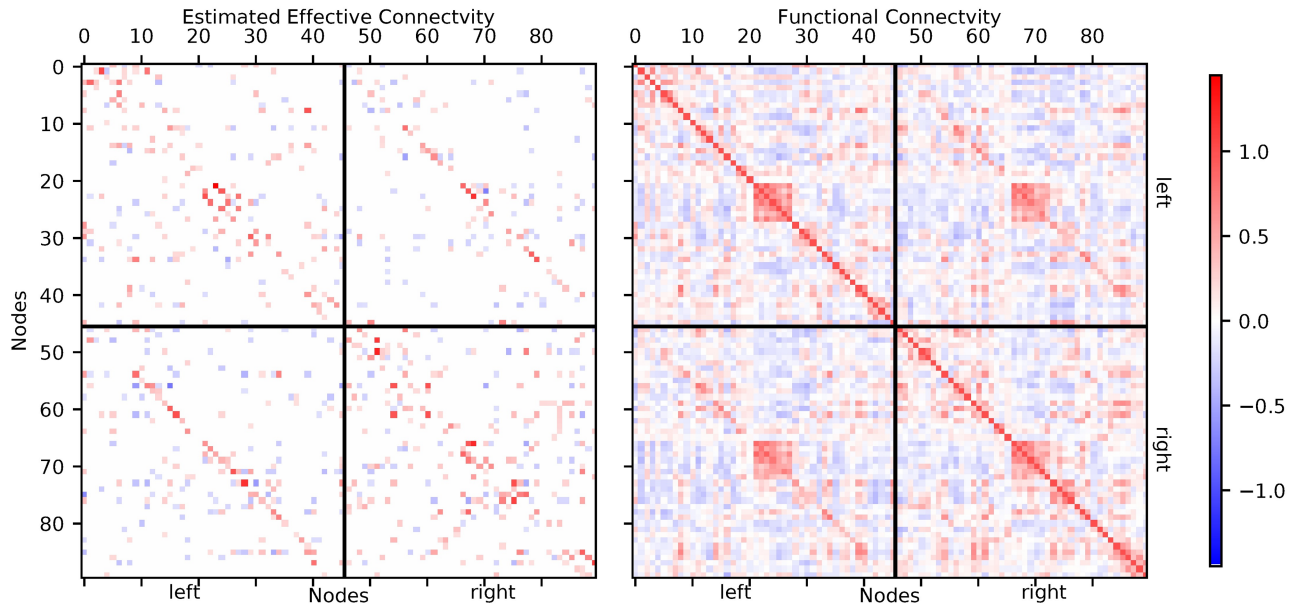


Fig 6. Left panel: Directed connectivity estimated with our new method from one sample MREG data set. Voxels were parceled using the AAL90-atlas. In the top-left block of the connectivity matrix connections within the left hemisphere are shown, in the lower-right block connections within the right hemisphere. The off-diagonal blocks represent the inter-hemispheric connections from the left to the right hemisphere (lower left) and from the right to the left hemisphere (top right). The strength of all connections is color coded, with red representing positive (excitatory) connections and blue representing negative (inhibitory) connections. Only the strongest 10% of connections are shown. Right panel: Functional connectivity matrix derived from the same data.

<https://doi.org/10.1371/journal.pcbi.1006056.g006>

the outcome of both methods, we only considered the existence of connections (binary and symmetric connectivity) and disregarded weights and directions (weighted nonsymmetric connectivity). One representative example of the comparison of both methods is shown in Fig 8. For RIC, 376 out of 810 possible connections were identified in four subjects or more out

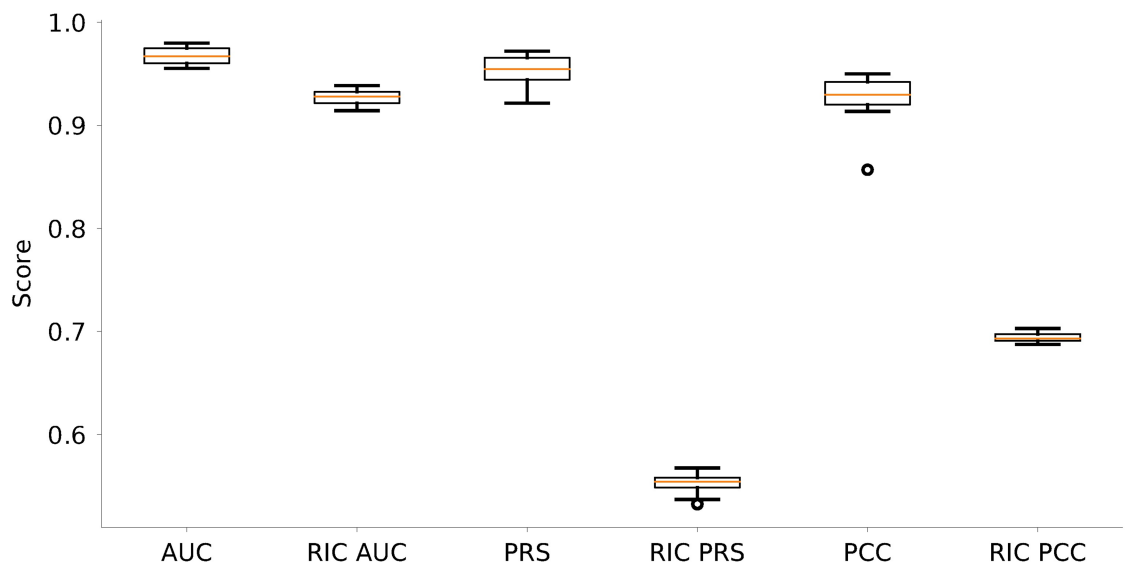


Fig 7. Comparison of performance with the Regularized Inverse Covariance (RIC) method based on numerical simulations of Ornstein-Uhlenbeck processes. Shown are the results from the reconstruction of 20 different networks with Erdős-Rényi connectivity profiles as described before (cf. Fig 2). AUC, PRS and PCC of our new method and of the RIC method, respectively, are shown side-by-side.

<https://doi.org/10.1371/journal.pcbi.1006056.g007>

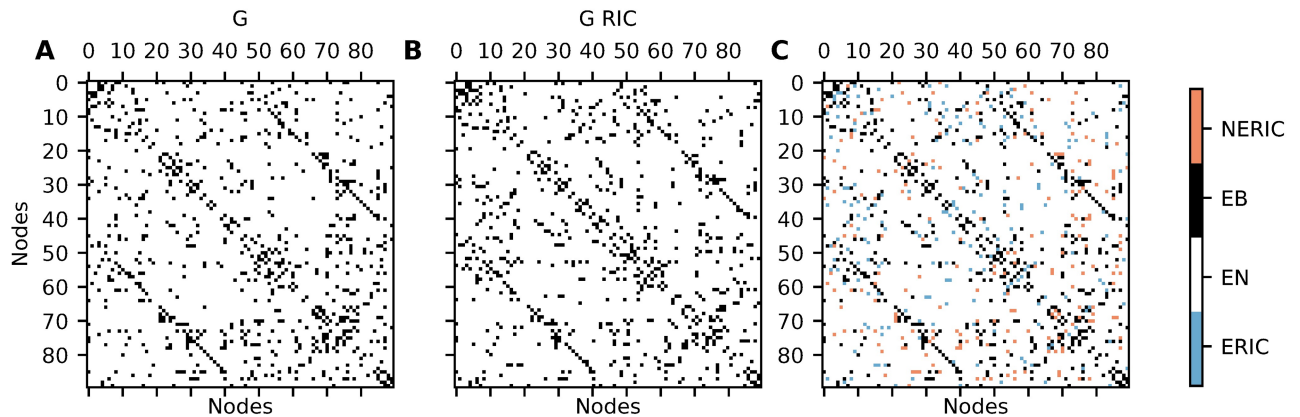


Fig 8. Estimated networks for one representative MREG dataset. The left panel shows the symmetrized network reconstructed with our estimation method, the middle panel shows the network found with the RIC method. The right panel shows the connections which are identified by both methods (EB, black), by none of the methods (EN, white), the connections found only by the RIC (ERIC, blue) and the connections found only by our method, but not by the RIC method (NERIC, red).

<https://doi.org/10.1371/journal.pcbi.1006056.g008>

of seven, the corresponding number for our method is 392 out of 810 possible connections. If any method produced directed networks with 10% connection probability at random, this would yield an average count of less than 25 connections (3% of 810 connections) that agree for least four out of seven independently generated networks. On average, 290.5 out of 4050 possible connections (undirected) are identified by both methods, 3530.5 connections were found by neither of the methods. This means that both methods agree on 3821 out of 4050 connections on average. The two methods disagreed on the remaining 229 connections.

Discussion

With the described method we can estimate directed and signed effective connectivity between neural populations from measured brain signals, based on zero-lag covariances only. To investigate the reliability of our estimated connections we used simulations of Ornstein-Uhlenbeck processes mimicking BOLD-related signals generated by interacting neuronal populations. Our method shows very good performance, if enough data is available and the observation noise is not too strong. Also, even in cases with relatively poor performance (e.g. if the network is too dense) more than 90% of the estimated connections have the correct sign. Applying the method on measured fast fMRI data, we found that about 34% of all identified connections have an inhibitory effect on their respective target population. In general, inhibitory synapses are mainly formed within local populations, and typically do not project to distant targets. An inhibitory connection between populations, however, can also be achieved by excitatory neurons preferentially terminating on the inhibitory neurons of the target region. The comparison with the Regularized Inverse Covariance (RIC) method shows good agreement with regard to the existence of connections. Directions cannot be disambiguated with the RIC method. Our results based on simulated surrogate data reflect what one would expect from the design of an estimation procedure. For large, sparse networks with sufficiently strong interaction, our network estimation procedure works reliably. However, as expected if the network is not sparse, or the time series is too short, the quality of the estimate drops. Nevertheless, in most cases the main interest lies on the strongest connections, which can be reliably estimated with our method even when the network is not sparse. For the experimental data shown, individual connections may be unreliable because of the limited size of the dataset. Also, it is unclear whether the biological network to be analyzed is really sparse, and if the assumption of

pairwise independent external input is really justified. On the other hand, due to the higher likelihood of a coupling between close-by areas and between inter-hemispheric counterparts, the resulting network looks plausible. For interpreting individual connections longer recordings would certainly be beneficial. Also, one could then use temporal information from additional frequency bands. Of high interest is also the comparison with structural measures as the ones obtained by diffusion tensor imaging. To the best knowledge of the authors, this is the first time that effective whole-brain connectivity has been estimated from zero-lag covariances. Other methods [8] rely on lagged covariances, where the correct lag parameter is critical, and needs to be inferred from the exponential decay of the observed auto-covariances. Also, our proposed method is the only one that can detect directed inhibitory connections on the whole-brain scale. The estimation procedure is fast and easy to apply. As it uses no temporal information, our method can also be applied on other data types that rely on the BOLD effect, e.g. fNIRS, but also data types measuring electrical population activity directly. This makes it a good candidate for, among other things, studying changing connectivity in neurodegenerative diseases, like Parkinson's or Alzheimer's.

Conclusion

With the presented method we can estimate directed effective connectivity on a whole-brain scale. Also we are able to detect whether connections are excitatory or inhibitory. The estimation is possible based on zero-lag covariances, but can also be applied to frequency-resolved cross spectral densities.

Supporting information

S1 Text. Supporting information.
(PDF)

Acknowledgments

We thank Uwe Grauer from the Bernstein Center Freiburg as well as Bernd Wiebelt and Michael Janczyk from the Freiburg University Computing Center for their assistance with HPC issues.

Author Contributions

Conceptualization: Jürgen Hennig, Pierre LeVan, Stefan Rotter.

Data curation: Carolin Lennartz, Pierre LeVan.

Formal analysis: Jonathan Schiefer, Alexander Niederbühl, Volker Pernice, Carolin Lennartz.

Funding acquisition: Jürgen Hennig, Stefan Rotter.

Investigation: Jonathan Schiefer, Alexander Niederbühl, Volker Pernice, Carolin Lennartz, Pierre LeVan.

Methodology: Volker Pernice.

Project administration: Stefan Rotter.

Resources: Jürgen Hennig, Stefan Rotter.

Software: Jonathan Schiefer, Alexander Niederbühl, Volker Pernice, Carolin Lennartz.

Supervision: Volker Pernice, Jürgen Hennig, Pierre LeVan, Stefan Rotter.

Validation: Jonathan Schiefer, Alexander Niederbühl, Volker Pernice.

Visualization: Jonathan Schiefer, Alexander Niederbühl.

Writing – original draft: Jonathan Schiefer.

Writing – review & editing: Jonathan Schiefer, Alexander Niederbühl, Volker Pernice, Carolin Lennartz, Pierre LeVan, Stefan Rotter.

References

1. Li B, Razi A, Friston KJ. Editorial: Mapping Psychopathology with fMRI and Effective Connectivity Analysis. *Front. Hum. Neurosci.* 2017; 11:151. <https://doi.org/10.3389/fnhum.2017.00151> PMID: 28408874
2. Friston KJ. Functional and effective connectivity: a review. *Brain connectivity.* 2011; 1(1):13–36. <https://doi.org/10.1089/brain.2011.0008> PMID: 22432952
3. Pernice V, Staude B, Cardanobile S, Rotter S. How structure determines correlations in neuronal networks. *PLoS Comput. Biol.* 2011; 7;. <https://doi.org/10.1371/journal.pcbi.1002059> PMID: 21625580
4. Havlicek M, Roebroek A, Friston KJ, Gardumi A, Ivanov D, Uludag K. On the importance of modeling fMRI transients when estimating effective connectivity: A dynamic causal modeling study using ASL data. *NeuroImage.* 2017; 155(July 2016):217–233. <https://doi.org/10.1016/j.neuroimage.2017.03.017> PMID: 28323165
5. Stephan KE, Friston KJ. Analyzing effective connectivity with fMRI. *Wiley interdisciplinary reviews Cognitive science.* 2010; 1(3):446–459. <https://doi.org/10.1002/wcs.58> PMID: 21209846
6. Smith JF, Pillai A, Chen K, Horwitz B. Identification and validation of effective connectivity networks in functional magnetic resonance imaging using switching linear dynamic systems. *NeuroImage.* 2010; 52(3):1027–1040. <https://doi.org/10.1016/j.neuroimage.2009.11.081> PMID: 19969092
7. Freestone DR, Karoly PJ, Nešić D, Aram P, Cook MJ, Grayden DB. Estimation of effective connectivity via data-driven neural modeling. *Front. Neurosci.* 2014; 8:383. <https://doi.org/10.3389/fnins.2014.00383> PMID: 25506315
8. Gilson M, Moreno-Bote R, Ponce-Alvarez A, Ritter P, Deco G. Estimation of Directed Effective Connectivity from fMRI Functional Connectivity Hints at Asymmetries of Cortical Connectome. *PLoS computational biology.* 2016. <https://doi.org/10.1371/journal.pcbi.1004762>
9. Ting CM, Seghouane AK, Member S, Salleh SH, Noor AM. Estimating Effective Connectivity from fMRI Data Using Factor-based Subspace Autoregressive Models. *IEEE Signal Processing Letters.* 2015; 22(6):757–761. <https://doi.org/10.1109/LSP.2014.2365634>
10. Roebroek A, Formisano E, Goebel R. The identification of interacting networks in the brain using fMRI: Model selection, causality and deconvolution. *NeuroImage.* 2011; 58(2):296–302. <https://doi.org/10.1016/j.neuroimage.2009.09.036> PMID: 19786106
11. Mehta-Pandjee G, Robinson PA, Henderson JA, Aquino KM, Sarkar S. Inference of direct and multi-step effective connectivities from functional connectivity of the brain and of relationships to cortical geometry. *Journal of Neuroscience Methods.* 2017; 283:42–54. <https://doi.org/10.1016/j.jneumeth.2017.03.014> PMID: 28342831
12. Marrelec G, Krainik A, Duffau H, Doyon J, Benali H. Partial correlation for functional brain interactivity investigation in functional MRI. *NeuroImage.* 2006; 32:228–237. <https://doi.org/10.1016/j.neuroimage.2005.12.057> PMID: 16777436
13. Timme M, Casadiego J. Revealing networks from dynamics: an introduction. *Journal of Physics A: Mathematical and Theoretical.* 2014; 47(34):343001. <https://doi.org/10.1088/1751-8113/47/34/343001>
14. Gates KM, Molenaar PCM. Group search algorithm recovers effective connectivity maps for individuals in homogeneous and heterogeneous samples. *NeuroImage.* 2012; 63(1):310–319. <https://doi.org/10.1016/j.neuroimage.2012.06.026> PMID: 22732562
15. Ramsey J, Glymour M, Sanchez-Romero, Glymour C. A million variables and more: the Fast Greedy Equivalence Search algorithm for learning high-dimensional graphical causal models, with an application to functional magnetic resonance images. *International Journal of Data Science and Analytics.* 2016. <https://doi.org/10.1007/s41060-016-0032-z> PMID: 28393106
16. Pernice V, Rotter S. Reconstruction of sparse connectivity in neural networks from spike train covariances. *Journal of Statistical Mechanics: Theory and Experiment.* 2013; 2013(03):P03008. <https://doi.org/10.1088/1742-5468/2013/03/P03008>
17. Pearl J. *Causality: Models, Reasoning, and Inference.* New York, NY, USA: Cambridge University Press; 2000.

18. Rebane G, Pearl J. The recovery of causal poly-trees from statistical data. *Proceedings of the Third Workshop on Uncertainty in AI*; 1987; 222–228
19. Candes EJ, Tao T. Decoding by Linear Programming. *IEEE Trans Inf Theor*. 2005; 51(12):4203–4215. <https://doi.org/10.1109/TIT.2005.858979>
20. Abrudan T, Eriksson J, Koivunen V. Conjugate Gradient Algorithm for Optimization Under Unitary Matrix Constraint. *Signal Processing*. 2009; 89:1704–1714. <https://doi.org/10.1016/j.sigpro.2009.03.015>
21. Wen Z, Yin W. A feasible method for optimization with orthogonality constraints. *Mathematical Programming*. 2013; 142(1):397–434. <https://doi.org/10.1007/s10107-012-0584-1>
22. Horn RA, Johnson CR. *Matrix Analysis*. Cambridge University Press, Cambridge, MA, 1985
23. Welvaert M, Rosseel Y. A review of fMRI simulation studies. *PLoS ONE*. 2014; 9(7):e101953. <https://doi.org/10.1371/journal.pone.0101953> PMID: 25048024
24. Gardiner CW. *Handbook of stochastic methods for physics, chemistry and the natural sciences*. vol. 13 of Springer Series in Synergetics. 3rd ed. Berlin: Springer-Verlag; 2004.
25. Gillespie DT. Exact numerical simulation of the Ornstein-Uhlenbeck process and its integral. *Phys Rev E*. 1996; 54:2084–2091. <https://doi.org/10.1103/PhysRevE.54.2084>
26. Friston KJ, Fletcher P, Josephs O, Holmes A, Rugg MD, Turner R. Event-Related fMRI: Characterizing Differential Responses. *NeuroImage*. 1998; 7(1):30–40. <https://doi.org/10.1006/nimg.1997.0306> PMID: 9500830
27. Glover GH. Deconvolution of Impulse Response in Event-Related BOLD fMRI. *NeuroImage*. 1999; 9(4):416–429. <https://doi.org/10.1006/nimg.1998.0419> PMID: 10191170
28. Assländer J, Zahneisen B, Hugger T, Reiser M, Lee HL, LeVan P, Hennig J. Single shot whole brain imaging using spherical stack of spirals trajectories. *NeuroImage*. 2013; 73:59–70. <https://doi.org/10.1016/j.neuroimage.2013.01.065>
29. LeVan P, Akin B, Hennig J. Fast imaging for mapping dynamic networks. *NeuroImage*. 2017 <https://doi.org/10.1016/j.neuroimage.2017.08.029>
30. Glover GH, Li TQ, Ress D. Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. *Magnetic Resonance in Medicine*. 2000; 44(1):162–167. [https://doi.org/10.1002/1522-2594\(200007\)44:1%3C162::AID-MRM23%3E3.0.CO;2-E](https://doi.org/10.1002/1522-2594(200007)44:1%3C162::AID-MRM23%3E3.0.CO;2-E) PMID: 10893535
31. Nie L, Yang X, Matthews PM, Xu ZW, and Guo YK. Inferring functional connectivity in fMRI using minimum partial correlation. *International Journal of Automation and Computing*, 14:371–385, 2017. <https://doi.org/10.1007/s11633-017-1084-9>
32. Ryali S, Chen T, Supekar K, Tu T, Kochalka J, Cai W and Menon V. Multivariate dynamical systems-based estimation of causal brain interactions in fMRI: Group-level validation using benchmark data, neurophysiological models and human connectome project data. *Journal of Neuroscience Methods*, 268:142–153, 2016. <https://doi.org/10.1016/j.jneumeth.2016.03.010> PMID: 27015792
33. Hyvärinen A and Smith SM. Pairwise likelihood ratios for estimation of non-gaussian structural equation models. *Journal of Machine Learning Research*, 14(1):111–152, 2013.
34. Smith SM, Miller KL, Salimi-Khorshidi G, Webster M, Beckmann CF, Nichols TE, Ramsey JD, Woolrich MW. Network modelling methods for fMRI. *NeuroImage* 2011; 54:2:875–891. <https://doi.org/10.1016/j.neuroimage.2010.08.063> PMID: 20817103
35. Zaghlool SB and Wyatt CL. Missing data estimation in fMRI dynamic causal modeling. *Frontiers in Neuroscience*, 8:191, 2014. <https://doi.org/10.3389/fnins.2014.00191> PMID: 25071435