

1 RESEARCH

2 **Whole-Brain Causal Discovery Using fMRI**

3 Fahimeh Arab¹, AmirEmad Ghassami², Hamidreza Jamalabadi³,
4 Megan A. K. Peters^{4,5,6}, and Erfan Nozari^{1,7,8}

5 ¹Department of Electrical and Computer Engineering, University of California, Riverside, CA

6 ²Department of Mathematics and Statistics, Boston University, USA

7 ³Department of Psychiatry and Psychotherapy, Phillips University of Marburg, Germany

8 ⁴Department of Cognitive Sciences, University of California, Irvine, USA

9 ⁵Center for the Neurobiology of Learning & Memory, University of California, Irvine, USA

10 ⁶Program in Brain, Mind, & Consciousness, Canadian Institute for Advanced Research, Canada

11 ⁷Department of Mechanical Engineering, University of California, Riverside, USA

12 ⁸Department of Bioengineering, University of California, Riverside, USA

13 **Keywords:** fMRI, Causal discovery, Brain networks, Statistical algorithms, Cognitive neuroscience

ABSTRACT

14 Despite significant research, discovering causal relationships from fMRI remains a challenge.
15 Popular methods such as Granger Causality and Dynamic Causal Modeling fall short in handling
16 contemporaneous effects and latent common causes. Methods from causal structure learning
17 literature can address these limitations but often scale poorly with network size and need
18 acyclicity. In this study, we first provide a taxonomy of existing methods and compare their
19 accuracy and efficiency on simulated fMRI from simple topologies. This analysis demonstrates a
20 pressing need for more accurate and scalable methods, motivating the design of Causal discovery

21 for Large-scale Low-resolution Time-series with Feedback (CaLLTiF). CaLLTiF is a
22 constraint-based method that uses conditional independence between contemporaneous and
23 lagged variables to extract causal relationships. On simulated fMRI from the macaque
24 connectome, CaLLTiF achieves significantly higher accuracy and scalability than all tested
25 alternatives. From resting-state human fMRI, CaLLTiF learns causal connectomes that are highly
26 consistent across individuals, show clear top-down flow of causal effect from attention and default
27 mode to sensorimotor networks, exhibit Euclidean distance-dependence in causal interactions,
28 and are highly dominated by contemporaneous effects. Overall, this work takes a major step in
29 enhancing causal discovery from whole-brain fMRI and defines a new standard for future
30 investigations.

AUTHOR SUMMARY

31 Discovering causal relationships from fMRI data is challenging due to contemporaneous effects
32 and latent causes. Popular methods like Granger Causality and Dynamic Causal Modeling
33 struggle with these issues, especially in large networks. To address this, we introduce CaLLTiF, a
34 scalable method that uses both lagged and contemporaneous variables to identify causal
35 relationships. CaLLTiF outperforms various existing techniques in accuracy and scalability on
36 simulated fMRI data. When applied to human resting-state fMRI, it reveals consistent and
37 biologically-plausible patterns across individuals, with a clear top-down causal flow from
38 attention and default mode networks to sensorimotor areas. Overall, this work advances the field
39 of causal discovery in large-scale fMRI studies.

INTRODUCTION

40 A major step in the global drive for understanding the brain (Adams et al., 2020; Amunts et al.,
41 2016; Jorgenson et al., 2015; Okano, Miyawaki, & Kasai, 2015; Poo et al., 2016) is to move
42 beyond correlations and understand the causal relationships among internal and external factors –
43 a process often referred to as *causal discovery* (Assaad, Devijver, & Gaussier, 2022; Glymour,
44 Zhang, & Spirtes, 2019). When possible, causal discovery can be greatly simplified by

45 intervening in one variable and observing the effect in others. However, such interventions are
46 often costly and/or infeasible, necessitating the significantly more challenging task of causal
47 discovery from purely observational data.

48 A particularly rich set of observational data for the brain comes from functional MRI
49 (fMRI) (Goense & Logothetis, 2008; Winder AT & PJ, 2017). The whole-brain coverage allowed
50 by fMRI is valuable for causal discovery not only because it allows for purely data-driven and
51 unbiased discovery of potentially unexpected causal relationships (Bressler & Menon, 2010;
52 Bullmore & Sporns, 2012; Fornito & Bullmore, 2015), but also because of the great extent to
53 which the presence of unobserved variables can complicate delineating causal adjacencies and
54 orientations (Entner & Hoyer, 2010; Gerhardus & Runge, 2020; Hasan, Hossain, & Gani, 2023;
55 Malinsky & Spirtes, 2018). Nevertheless, many characteristics of fMRI also make causal
56 discovery challenging, including its large dimensionality, low temporal resolution, and indirect
57 reflection of underlying neural processes (Ramsey et al., 2010).

58 This has motivated a large and growing body of literature on causal discovery from fMRI. A
59 common approach for causal discovery using neuroimaging and neurophysiology data is Granger
60 Causality (GC) (Seth, Barrett, & Barnett, 2015). GC has a long history in neuroscience (Barnett
61 & Seth, 2014; Seth et al., 2015), but also has well-known limitations, including its lack of ability
62 to account for contemporaneous causal relationships and the presence of latent nodes (see
63 Supplementary Note 1 for a formal definition of ‘contemporaneous’ causal effects). The former is
64 particularly important for fMRI. The temporal resolution in fMRI is typically within a few
65 hundred milliseconds to several seconds (Huettel, Song, & McCarthy, 2009), which is about one
66 order of magnitude slower than the time that it takes for neural signals to travel across the brain
67 (Nunez & Srinivasan, 2006; Sutton & Begleiter, 1979; Sutton, Braren, Zubin, & John, 1965).
68 Therefore, from one fMRI sample to the next, there is enough time for causal effects to flow
69 between almost all pairs of nodes in the network (cf. a related in-depth discussion in (Nozari,
70 Pasqualetti, & Cortés, 2019, Appendix A)). Such fast sub-TR interactions demonstrate
71 themselves as causal effects that appear to be “contemporaneous” and can even be cyclic, making
72 causal discovery significantly more challenging (cf. Supplementary Note 1). Similar to GC,
73 Dynamic Causal Modeling (DCM) has also been widely used with fMRI data (K. Friston,

⁷⁴ Harrison, & Penny, 2003; K. Friston et al., 2019; K. J. Friston, Kahan, Biswal, & Razi, 2014;
⁷⁵ Stephan & Roebroeck, 2012) and fundamentally relies on the temporal order of a generative
⁷⁶ dynamical model to infer causation from correlations, making it similarly unable to account for
⁷⁷ contemporaneous causal relationships (K. Friston, Moran, & Seth, 2013; K. J. Friston, 2011;
⁷⁸ Logothetis, 2008).

⁷⁹ Discovering causal relationships without reliance on time has been the subject of extensive
⁸⁰ research in the causal inference literature (Glymour et al., 2019; Pearl, 1988, 2009a, 2009b;
⁸¹ Spirtes, Glymour, & Scheines, 2000; Spirtes & Zhang, 2016). A wide range of algorithmic
⁸² solutions have been proposed (Chickering, 2002a; Glymour et al., 2019; Henry & Gates, 2017;
⁸³ Meek, 1995, 1997; Pearl, 2009a; Ramsey et al., 2010; Shimizu, Hoyer, Hyvärinen, & Kerminen,
⁸⁴ 2006; Smith et al., 2011; Spirtes & Glymour, 1991; Spirtes, Meek, & Richardson, 1995), which
⁸⁵ are often classified based on their methodology into constraint-based (Dawid, 1979; Pearl, 1988,
⁸⁶ 2009b), noise-based (Shimizu, 2014; Shimizu et al., 2006), and score-based (Chickering, 2002b;
⁸⁷ Heckerman, Geiger, & Chickering, 1995). However, it has remained largely unknown which of
⁸⁸ these algorithms are best suited for whole-brain fMRI causal discovery and how they perform
⁸⁹ relative to one another in terms of accuracy and scalability.

⁹⁰ In this study, we first discuss and compare existing causal discovery algorithms for their
⁹¹ suitability for whole-brain fMRI, demonstrate a large gap between what causal discovery for
⁹² fMRI needs and what existing algorithms can achieve, propose CaLLTiF to address this gap, and
⁹³ demonstrate its higher accuracy and scalability on synthetic and real fMRI. Unless otherwise
⁹⁴ noted, all references to the words ‘graph’/‘network’ and ‘cycle’ mean a directed graph and
⁹⁵ directed cycle, respectively.

RESULTS

⁹⁶ *A Taxonomy of Causal Discovery for Whole-Brain fMRI*

⁹⁷ A vast array of algorithmic solutions exist for learning causal graphs from observational data, but
⁹⁸ not all are suitable for fMRI data. We selected a subset of state-of-the-art algorithms suitable for
⁹⁹ whole-brain fMRI data based on four criteria: (1) ability to learn cycles, (2) ability to learn
¹⁰⁰ contemporaneous effects, (3) assuming complete coverage of relevant variables in observed data,

Table 1: List of causal discovery methods suitable for use with whole-brain fMRI, divided by methodological category (constraint-, noise-, and score-based). All these methods (1) allow for cycles, (2) allow for contemporaneous effects, (3) assume complete coverage of relevant variables in observed data, and (4) learn linear relationships. The FASK algorithm is fundamentally hybrid and therefore listed as both constraint-based and noise-based.

Category Type \ Category	Constraint-based	Noise-based	Score-based
Time-series	PCMCI (Runge, Nowack, Kretschmer, Flaxman, & Sejdinovic, 2019), PCMCI ⁺ (Runge, 2020)	VARLiNGAM (Hyvärinen, Zhang, Shimizu, & Hoyer, 2010)	DYNOTEARs (Pamfil et al., 2020)
Cross-sectional with cycles	FASK (Sanchez-Romero et al., 2019)	FASK (Sanchez-Romero et al., 2019), LiNG (Lacerda, Spirtes, Ramsey, & Hoyer, 2008)	DGlearn (Ghassami, Yang, Kiyavash, & Zhang, 2020)

and (4) linearity (see Discussions). Table 1 shows several state-of-the-art methods that satisfy criteria (1)-(4). Multivariate Granger Causality (MVGC) (Barnett & Seth, 2014; Granger, 1969) does not satisfy criteria (2), but we still included it in our subsequent analyses due to its popularity in neuroscience (Ding, Chen, & Bressler, 2006; Goebel, Roebroeck, Kim, & Formisano, 2003; Liao et al., 2010; Roebroeck, Formisano, & Goebel, 2005). On the other hand, we excluded LiNG (Ramsey et al., 2018) from further analysis since it is considered by its proposers as generally inferior to the hybrid FASK algorithm (Sanchez-Romero et al., 2019). We also chose FASK for implementation over GANGO (Rawls, Kummerfeld, Mueller, Ma, & Zilverstand, 2022), a similarly hybrid method with the additional caveat of not having a unified publicly available code distribution.

We compared the accuracy of the resulting list of algorithms (MVGC, PCMCI, PCMCI⁺, VARLiNGAM, DYNOTEARs, FASK, and DGlearn) using simulated fMRI data from a benchmark of simple (5-10 nodes) networks introduced in (Sanchez-Romero et al., 2019). The ground truth graphs are shown in Figure 1a, and details on the fMRI time series generation for each node in these graphs are provided under Methods. To evaluate the success of each algorithm, we treated the causal discovery problem as a binary classification problem for each directed edge

117 and calculated the resulting F1 score, both for the directed graphs as well as their underlying
118 undirected graphs (see Methods for details). Figure 1b illustrates the distribution of F1 scores for
119 all algorithms, combined across nine simple networks. The results show that the PCMCI
120 algorithm achieved significantly higher median F1 score compared to all other algorithms over
121 the directed graphs (all Cohen's $d > 0.23$ and $p < 10^{-4}$, pairwise one-sided Wilcoxon
122 signed-rank test) and compared to all but DYNOTEARs over the underlying undirected graphs
123 (Supplementary Figure 3, all Cohen's $d > 0.44$ and $p < 10^{-29}$, pairwise one-sided Wilcoxon
124 signed-rank test. Also see Supplementary Figures 4 and 5 for precision and recall). The PCMCI
125 algorithm also has the smallest computational complexity on simple networks, as seen from
126 Figure 1c. Furthermore, our results indicate that FASK, DGlearn, and PCMCI⁺ (at their best
127 values of hyperparameters) do not scale well with network size, forcing us to exclude them from
128 further analysis as we move on to larger networks (see Supplementary Figures 13, 14, and 15).

129 Next, we compared the remaining four algorithms (PCMCI, MVGC, DYNOTEARs, and
130 VARLiNGAM) on a larger-scale, more realistic simulated benchmark. The graph shown in
131 Figure 2a, referred to as the Macaque_SmallDegree network, derived from the original
132 Macaque_Full anatomical connectome and pruned to achieve average in-degree and out-degree of
133 1.8, consists of 28 nodes and 52 directed edges (Sanchez-Romero et al., 2019) but the generative
134 model used to simulate fMRI data from this graph remains the same (see Methods for details).
135 The distributions of F1 scores are shown in Figure 2b. PCMCI and MVGC achieved very similar
136 success in learning both the directed graph and its underlying undirected graph, while
137 significantly outperforming DYNOTEARs and VARLiNGAM. A similar result is obtained when
138 comparing adjacency F1 scores for detecting the network's underlying undirected graph
139 (Supplementary Figure 17, also see Supplementary Figures 18 and 19 for precision and recall).
140 As far as execution time is concerned, however, MVGC showed a significant advantage over
141 PCMCI (Figure 2c). Therefore, despite its simplistic nature, MVGC was found most successful in
142 causal discovery from Macaque_SmallDegree fMRI data (but also see Figure 3).

143 **CaLLTiF: A New Causal Discovery Algorithm for Whole-Brain fMRI**

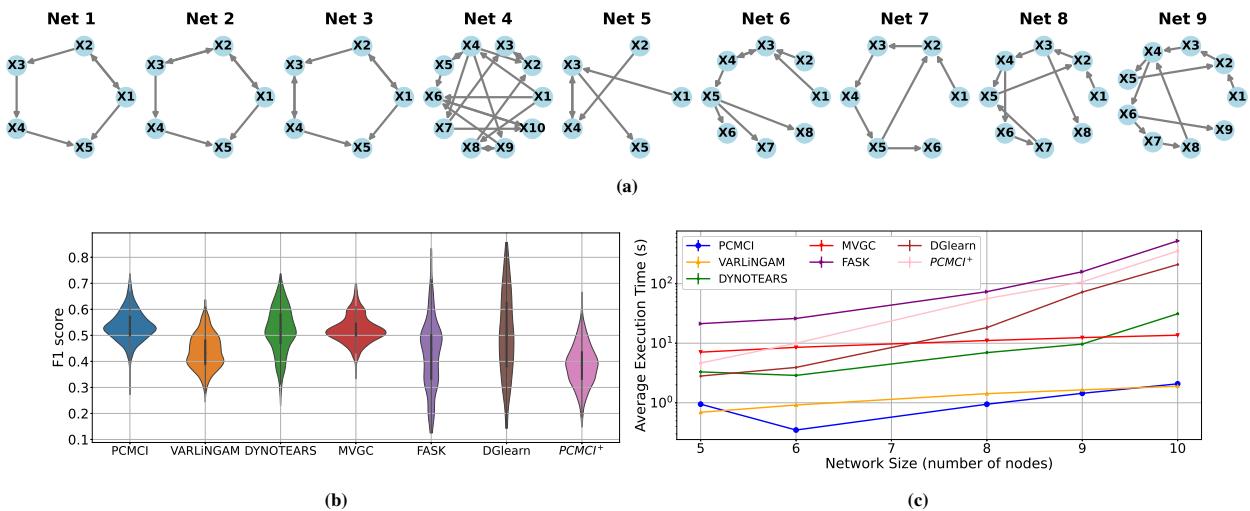


Figure 1: Results of comparing several state-of-the-art causal discovery algorithms over a benchmark of simulated fMRI (Sanchez-Romero et al., 2019) generated from simple networks with 5-10 nodes each. (a) Ground truth graphs of the simple networks in the benchmark. Despite all being small-scale, the graphs vary widely in their density, number of cycles, etc. (b) F1 scores of seven state-of-the-art algorithms (six from Table 1 and MVGC) for correctly identifying the full (directed) graphs. All methods are evaluated using optimized values of their respective hyperparameters (see Methods). The benchmark data includes 60 repetitions of fMRI data from each of the 9 graphs, so each violin plot is based on 540 F1 score samples. The PCMCI method achieves the highest median F1 score, both directed and undirected (see Supplementary Figure 3). (c) The mean execution time (averaged over all 60 repetitions) of different algorithms as a function of network size. Note the logarithmic scaling of the vertical axis. Error bars, though hardly visible, show 1 s.e.m.

144 The best-performing algorithms on Macaque_SmallDegree, i.e., PCMCI and MVGC, suffer from
 145 three main drawbacks: (1) poor scalability (only for PCMCI), (2) inability to learn directed
 146 contemporaneous effects (PCMCI only learns undirected contemporaneous effects while MVGC
 147 learns none), and (3) having sparsity-controlling hyperparameters that are subjectively selected in
 148 the absence of ground-truth graphs. In this section, we describe the design of a new algorithm
 149 based on PCMCI that mitigates these drawbacks and demonstrate its superior performance over
 150 existing methods.

151 Our first modification to PCMCI is with regard to scalability and computational complexity.
 152 The computational complexity of PCMCI depends heavily on the value of its ‘PC Alpha’
 153 hyperparameter, which controls the sparsity of the set of potential common causes on which the

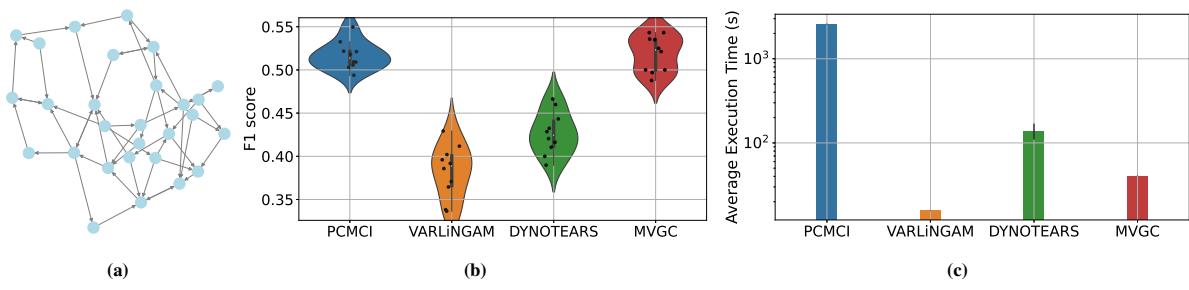


Figure 2: Comparing the scalable subset of algorithms from Figure 1 over simulated fMRI data from the Macaque_SmallDegree benchmark (Sanchez-Romero et al., 2019). (a) ground-truth Macaque_SmallDegree network. (b) F1 scores of identifying the directed graph. Each distribution consists of 10 F1 scores calculated based on 10 repetitions of simulated data from the same underlying graph. (c) Mean execution times for each method (error bars show one standard deviation).

154 algorithm conditions when checking the conditional independence of each pair of nodes
 155 (Supplementary Figures 16a and 37). Higher values of PC Alpha make these sets denser and
 156 accordingly decrease statistical power in the subsequent conditional independence tests,
 157 *ultimately conditioning on all other nodes (and all of their lags) when PC Alpha = 1.*
 158 Nevertheless, interestingly, our experiments on the Macaque_SmallDegree data show that the
 159 maximum achievable accuracy of PCMCI (i.e., F1 score maximized over Alpha Level for each
 160 fixed value of PC Alpha) monotonically increases with PC Alpha, reaching its maximum at PC
 161 Alpha = 1 (Supplementary Figures 16b and 16c). Therefore, while this may seem to cause a
 162 trade-off between accuracy and scalability, it is in fact an opportunity for maximizing both. At PC
 163 Alpha = 1, the PC part of PCMCI (a.k.a. the S1 algorithm in (Runge et al., 2019)) is theoretically
 164 guaranteed to return a complete conditioning set for all pairs of nodes, and can thus be skipped
 165 entirely. The PC part is further responsible for the poor scalability of PCMCI. Thus its removal
 166 significantly improves the computational efficiency of the resulting algorithm *without*
 167 *compromising accuracy* (cf. Discussions for a potential explanation of why conditioning on all
 168 other nodes may improve accuracy despite lowering statistical power).

169 Our second modification addresses the lack of directed contemporaneous causal effects (see
 170 Introduction for why these effects are particularly important in fMRI). By default, MVGC returns
 171 no contemporaneous edges and PCMCI returns $\circ-\circ$ ones which only indicate the presence of

172 significant partial correlations but does not resolve between \rightarrow , \leftarrow , or \leftrightharpoons . However, we know from
173 decades of tract tracing studies that reciprocal connections are significantly more common than
174 unidirectional connections in the primate brain (Felleman & Van Essen, 1991; Markov et al.,
175 2014; Tigges, Spatz, & Tigges, 1973). Therefore, we replace all $\circ\text{---}\circ$ edges returned by PCMCI
176 by the more likely choice of \leftrightharpoons . The only exception comes from (the often minority of) pairs of
177 nodes that have a *lagged* directed edge between them (i.e., an edge of the form
178 $X_i(t - \tau) \rightarrow X_j(t)$ for $\tau > 0$, see Methods), in which case we leave the direction of the
179 contemporaneous effect between them the same as their lagged effect.

180 Figure 3 shows how the resulting CaLLTiF algorithm performs on a synthetic fMRI dataset
181 generated from the significantly larger macaque structural connectome called Macaque_Full with
182 91 nodes and 1,615 ground-truth edges (Figure 3a, see also Methods). CaLLTiF has a
183 significantly higher F1 score compared to PCMCI, VARLiNGAM, DYNOTEARs, and MVGC in
184 its discovery of the directed graph (Figure 3b, all Cohen's $d > 15$, $p < 10^{-6}$, one-sided Wilcoxon
185 signed-rank test) as well as its underlying undirected graph (Figure 3c, all Cohen's $d > 7$,
186 $p < 10^{-6}$, one-sided Wilcoxon signed-rank test). Precisions and recalls are also shown in
187 Supplementary Figures 26 and 27. We also compared CaLLTiF (and PCMCI) against a
188 middle-ground ‘Mixed-PCMCI’ variant where the $\circ\text{---}\circ$ edges returned by PCMCI are used only
189 in the computation of adjacency F1 score (Supplementary Figures 28-30, see also Methods).
190 Mixed-PCMCI benefits from contemporaneous effects as much as CaLLTiF in terms of adjacency
191 F1 score, but not so in terms of full F1 score, further motivating the inclusion of directed
192 contemporaneous connections as done in CaLLTiF. Detailed performances of all compared
193 algorithms are provided in Supplementary Figures 29-36.

194 Finally, the third aspect in which CaLLTiF departs from PCMCI is the choice of
195 sparsity-controlling hyperparameter ‘Alpha Level’. Most, if not all, algorithms for causal
196 discovery have at least one hyperparameter (often a threshold) that controls the sparsity of the
197 resulting graphs. Different from PC Alpha described earlier, Alpha Level in PCMCI is the
198 standard type-I error bound in determining statistical significance in *each* partial correlation test
199 (cf. Supplementary Figure 16). By default, Alpha Level is selected subjectively, based on domain
200 knowledge and expected level of sparsity. However, in CaLLTiF, we select Alpha Level

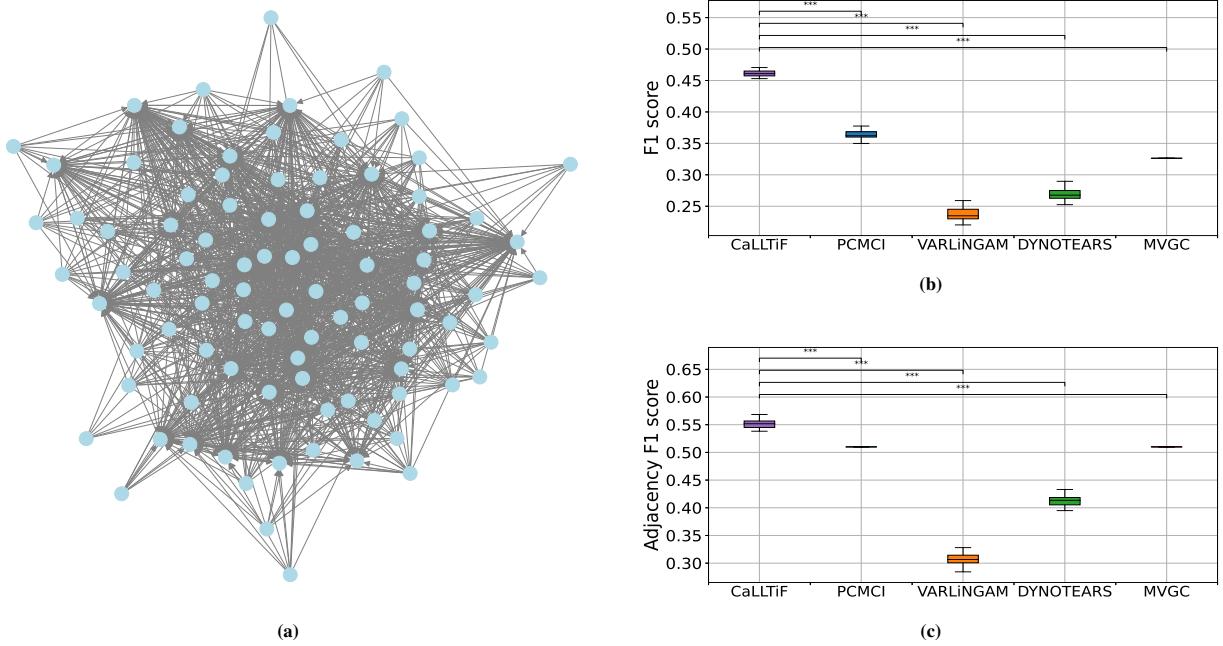


Figure 3: Comparisons between the proposed algorithm (CaLLTiF) and state-of-the-art alternatives over simulated fMRI from the Macaque_Full connectome. (a) ground-truth Macaque_Full network (See Supplementary Figure 25 for the heat map of the directed connectivity matrix). (b,c) Distributions of F1 scores for CaLLTiF and state-of-the-art alternatives in discovering the directed graph (b) and its underlying undirected graph (c). For all repetitions, the best performance of MVGC occurs at $\alpha = 0.5$ which returns a complete graph, hence the point distributions for MVGC. *** denotes $p < 0.001$. All statistical comparisons are performed using a one-sided Wilcoxon signed-rank test. In all boxplots, the center line represents the median, the box spans the interquartile range (IQR), and the whiskers extend up to 1.5 times the IQR from the box limits.

201 objectively based on a novel method for correction for multiple comparisons (see Methods) that
202 occur when collapsing a time-series graph over lagged variables into a final summary graph. This
203 step is critical, particularly in the absence of ground-truth connectivity, to ensure that we have
204 statistical confidence in every edge of the final summary graph returned by CaLLTiF.

205 In summary, CaLLTiF starts by constructing an extended time-lagged graph among all the
206 variables $X_i(t - \tau), i = 1, \dots, n$ and all lags $\tau = 0, 1, \dots, \tau_{\max}$. To establish a causal link
207 between any pair of variables $X_i(t - \tau)$ and $X_j(t)$, CaLLTiF performs a conditional
208 independence test (using linear partial correlation) between $X_i(t - \tau)$ and $X_j(t)$, conditioned on
209 all other lagged variables ($X_k(t - s), s = 1, \dots, \tau_{\max}$). A causal link is established if the null
210 hypothesis of conditional independence is rejected at a significance threshold ‘Alpha Level’. By
211 default, ‘Alpha Level’ is selected based on CaLLTiF’s type I error control, but it can also be
212 optimized in simulated data using ground-truth knowledge. If $\tau > 0$, the direction of the edge is
213 clearly $X_i(t - \tau) \rightarrow X_j(t)$. When $\tau = 0$, CaLLTiF returns $X_i(t) \leftarrow X_j(t)$ if no other edges exist
214 between X_i and X_j at higher lags, and places the edge(s) consistent with the corresponding
215 lagged direction(s) otherwise. Finally, the extended time-lagged graph is collapsed into a
216 summary graph by taking an OR operation for each edge across all lags. For further details about
217 CaLLTiF, see Methods. A pseudocode for CaLLTiF is given in Algorithm 1 and a formal analysis
218 of its computational complexity can be found in Supplementary Note 4.

219 ***Causal Discovery from Resting-State Human fMRI***

220 We next applied CaLLTiF on resting-state fMRI from 200 subjects from the Human Connectome
221 Project (HCP) (See Methods). Each scan from each subject was parcellated into 100 cortical and
222 16 subcortical regions. CaLLTiF was then performed on all four resting-state scans for each
223 subject, resulting in one causal graph per individual.

224 **Learned causal graphs are highly consistent across subjects.** Despite individual differences,
225 a remarkably common causal connectome emerged across subjects. Figure 4a shows the average
226 causal graph among the subjects and Figure 4b shows the intersection graph that contained the
227 edges *common across all subjects*. Due to the binary nature of individual graphs, the former can
228 also be viewed as a matrix of probabilities, where entry (i, j) shows the probability of region i

229 causing region j across all subjects. As a result of the significant commonalities that exist in the
230 causal graphs among subjects, the average graph has a bimodal distribution, with the vast
231 majority of average weights being close to either 0 or 1. These extreme values of average weights
232 can also be seen as a measure of the confidence of the algorithm in the presence or lack of most
233 edges, and have a clear contrast ($p = 0$, Kolmogorov-Smirnov test) with the weights of the
234 average of randomized surrogate graphs generated *independently* across subjects (Supplementary
235 Figure 44). In the absence of a ground truth causal connectivity for direct comparison, such
236 strong commonalities among subjects serve as an alternative measure of validation and provide
237 insights into the general patterns and characteristics of the causal relationships in a resting brain.

238 Nodal centralities also show strong consistency among the subjects. Figure 4c shows the nodal
239 degrees for all subjects (gray lines) as well as its mean across subjects (black line, also depicted in
240 Figure 4d). Statistically significant differences exist between the degree distributions of many
241 pairs of nodes (about 90% of the pairs have $p < 0.001$, two-sided Wilcoxon signed-rank test,
242 computed between nodal degrees of each pair of parcels), while significant correlations exist
243 between nodal degrees for all pairs of subjects (all pairs have $0.56 \leq r \leq 0.96$, $p < 10^{-10}$,
244 Pearson Correlation Coefficient, computed based on the nodal degrees of each pair of subjects
245 separately). Note that the correlations of degrees across subjects are also remarkably higher than
246 what would be implied by the correlations of the graphs themselves (Supplementary Figure 43e).
247 Similar consistency for in-degree, out-degree, betweenness, and eigenvector centralities can be
248 observed among subjects (Supplementary Figures 45-48). Consistently across subjects, medial
249 ventral attention regions, cingulate cortices, and lateral primary sensorimotor areas show
250 particularly low nodal degrees across both hemispheres, whereas bilateral default mode areas,
251 particularly the left ventromedial prefrontal cortex, show notably strong nodal degrees. Bilateral
252 anterior thalami are particularly causally connected compared to other subcortical regions, even
253 though subcortical areas have significantly lower degrees than cortical areas in general, with
254 bilateral posterior thalami, nuclei accumbens, and globus pallidi showing the least causal
255 connections across the whole brain at rest.

256 Causal graphs are also sparser and more consistent across subjects compared to functional
257 connectivity. A major motivation for building causal connectomes is the removal of spurious

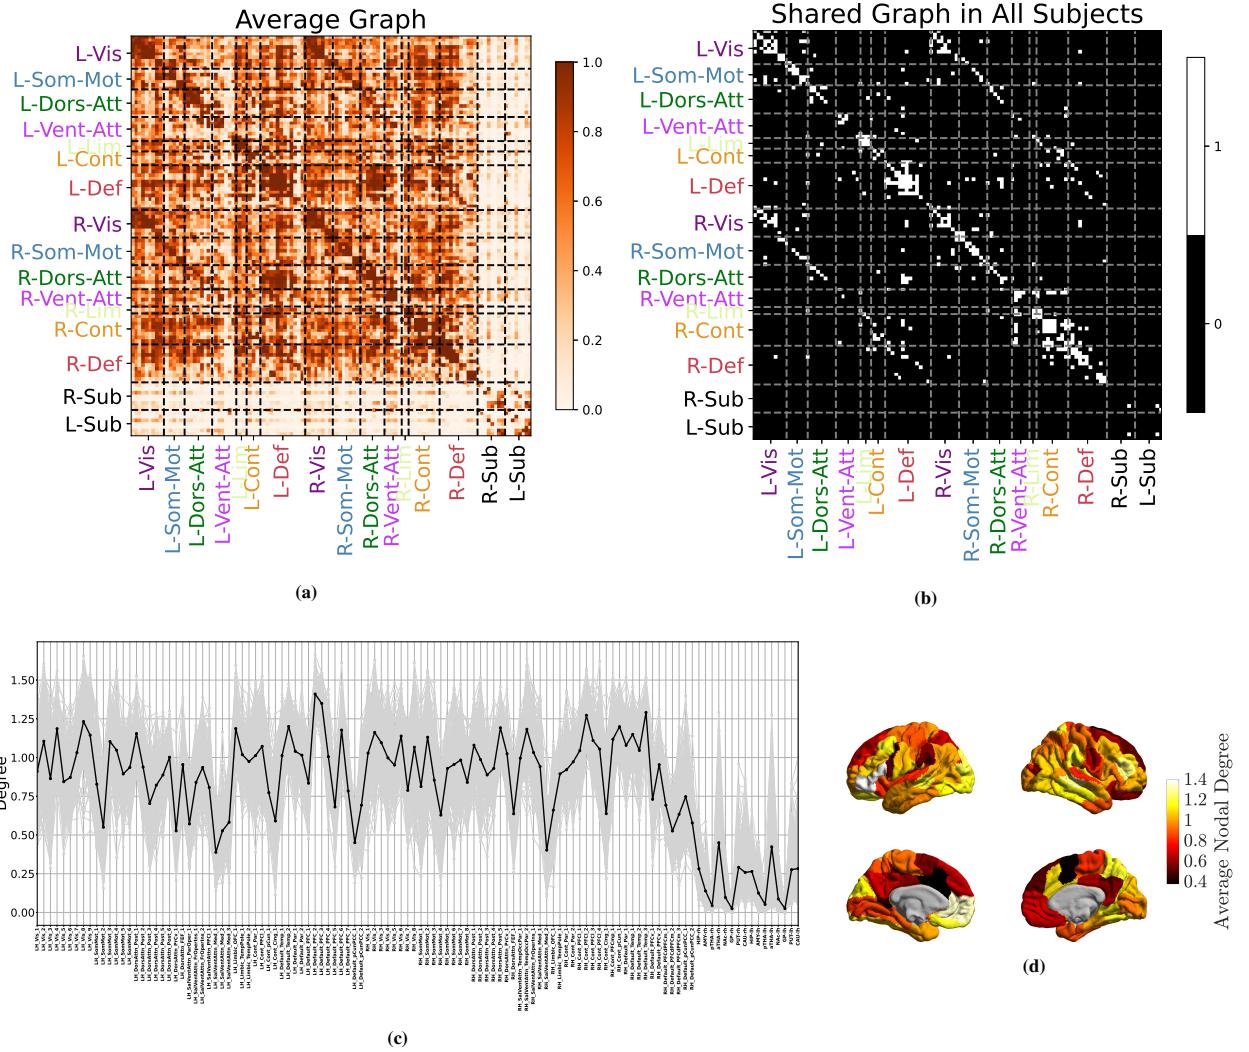


Figure 4: Learned causal graphs are highly consistent across subjects. (a) Average causal graph, computed as the mean of all the causal graphs of 200 subjects. The color of the (i, j) entry in this matrix shows how likely it is for node i to cause node j . A clear distinction can be seen in the causal links among cortical regions (colored labels) and subcortical ones (black labels). (b) The shared causal graph, containing only the edges that are present in all 200 subjects. This graph is dominated by diagonal elements (self- and within-subnetwork causation) and links among symmetrically located regions across the two hemispheres. (c) Distribution of nodal degree, computed separately for each node in the causal graph and each subject. Gray colors correspond to different subjects and the black line shows the average nodal degrees across subjects. (d) Average cortical nodal degree (black line in (c)) shown over the brain cortical surface.

258 connections in functional connectivity (FC) profiles that reflect mere correlation but no causation.
259 For causal graphs learned by CaLLTiF, we indeed observed significantly lower edge density
260 compared to FC graphs (see Methods for details on the computation of FC graphs)
261 (Supplementary Figures 43a,43b, no overlap existing between the support of the two
262 distributions). In fact, FC graphs included approximately 95% of CaLLTiF's discovered causal
263 edges (Supplementary Figure 43c), while only about half of all functional connectivity edges are
264 also causal (Supplementary Figure 43d). Interestingly, among the approximately 5% of causal
265 edges that were not in the FC graphs, the majority came from non-zero lags. This is remarkable,
266 given that causal edges from non-zero lags are significantly fewer in general (cf. Figure 6a), but
267 are fundamentally not discoverable by FC which only measures contemporaneous co-fluctuations.
268 Moreover, causal connectomes are significantly more consistent across subjects compared to FC
269 connectomes (Supplementary Figure 43e, Cohen'd > 2, $p < 0.001$, one-sided Wilcoxon
270 signed-rank test), further reinforcing the expectation that causal edges are “pruned” and more
271 reliable compared to functional edges.

272 **Net resting-state causal effect flows from attention and default mode to sensorimotor**
273 **networks.** One of the main advantages of directed causal connectomes over undirected functional
274 and structural connectomes is the former's ability to show the directed flow of causal effect
275 between brain regions. In graphs learned by CaLLTiF, nodal causal flows (outflow minus inflow,
276 see Methods) are also highly consistent across subjects (Figure 5c,5d), even though the two
277 notions of centrality are generally dissociated across parcels (Figure 5b and Supplementary
278 Figure 53). On average across all subjects, we observed particularly high causal flows
279 (source-ness) in several regions of bilateral medial ventral attention networks, specific dorsal
280 attention areas (ventral precentral, ventral frontal cortices, and frontal eye fields), and bilateral
281 hippocampi, even though subcortical areas are much less connected to the rest of the network in
282 general. In contrast, bilateral visual areas show the strongest negative causal flow (sink-ness)
283 across all subjects. There is also notable variability among parcels within a subnetwork, such as
284 the notable bilateral contrast between the strongly positive and weakly negative causal flows of
285 frontal and posterior parts of the dorsal attention network, respectively.

286 To better assess the overall net causal effects between different functional networks, we
287 computed an average “subnetwork graph” in which each node represents a functional cortical
288 network and edges denote thresholded average directed connectivity from one functional network
289 to another (see Methods for detailed computations). The result is illustrated in Figure 5a. Ventral
290 attention and visual networks are clearly the strongest source and sink of causal flow, respectively.
291 The dorsal attention and somatomotor networks are also a clear (though weaker) source and sink,
292 respectively. The default mode network (DMN) is also a net source of causal flow, even though its
293 outflows and inflows are nearly balanced. Similarly, the control and limbic networks have
294 near-zero causal flows (near-balanced inflow and outflow). Several directed paths, however, can
295 be seen from both attention and default mode networks to sensorimotor networks through the
296 limbic and control networks. Therefore, in summary, causal graphs learned by CaLLTiF show the
297 strongest net resting state causal effect to flow from the ventral and dorsal attention as well as the
298 default mode networks, through control and limbic networks, towards sensorimotor networks.
299 The DMN, control, and limbic networks have large average degrees (Supplementary Figure 52)
300 and near-balanced causal flows, making them hubs that largely distribute the flow of causal effect
301 in the resting-state causal connectome (see the Discussion section for a more detailed analysis of
302 this network).

303 **Casual graphs are strongly dominated by contemporaneous and lag-1 connections.** Given
304 that the final causal graph returned by CaLLTiF is a union over subgraphs at different lags (cf.
305 Methods), we can go back and ask how much causal effects in each lag have contributed to the
306 final graph. Figure 6a shows the percentage of edges in the final graph that exist *only* in one lag
307 (including lag 0, or contemporaneous edges). Increasing the lag order resulted in significantly
308 sparser single-lag subgraphs, which contributed less to the end result. In particular, approximately
309 70% of the end graphs came from lag 0 alone, a pattern that appears consistently across all
310 subjects (Supplementary Figure 57). Even further, such contemporaneous edges are substantially
311 stronger than edges from lags 1-3 (Figure 6b). This further confirms that the contemporaneous
312 effects are particularly important for fMRI, where most neural dynamics occur at timescales
313 shorter than 1 TR (typically shorter than 1-2 seconds). This is even the case in HCP data, with TR
314 = 0.72s which is among the shortest TRs currently available in fMRI research. That being said, all

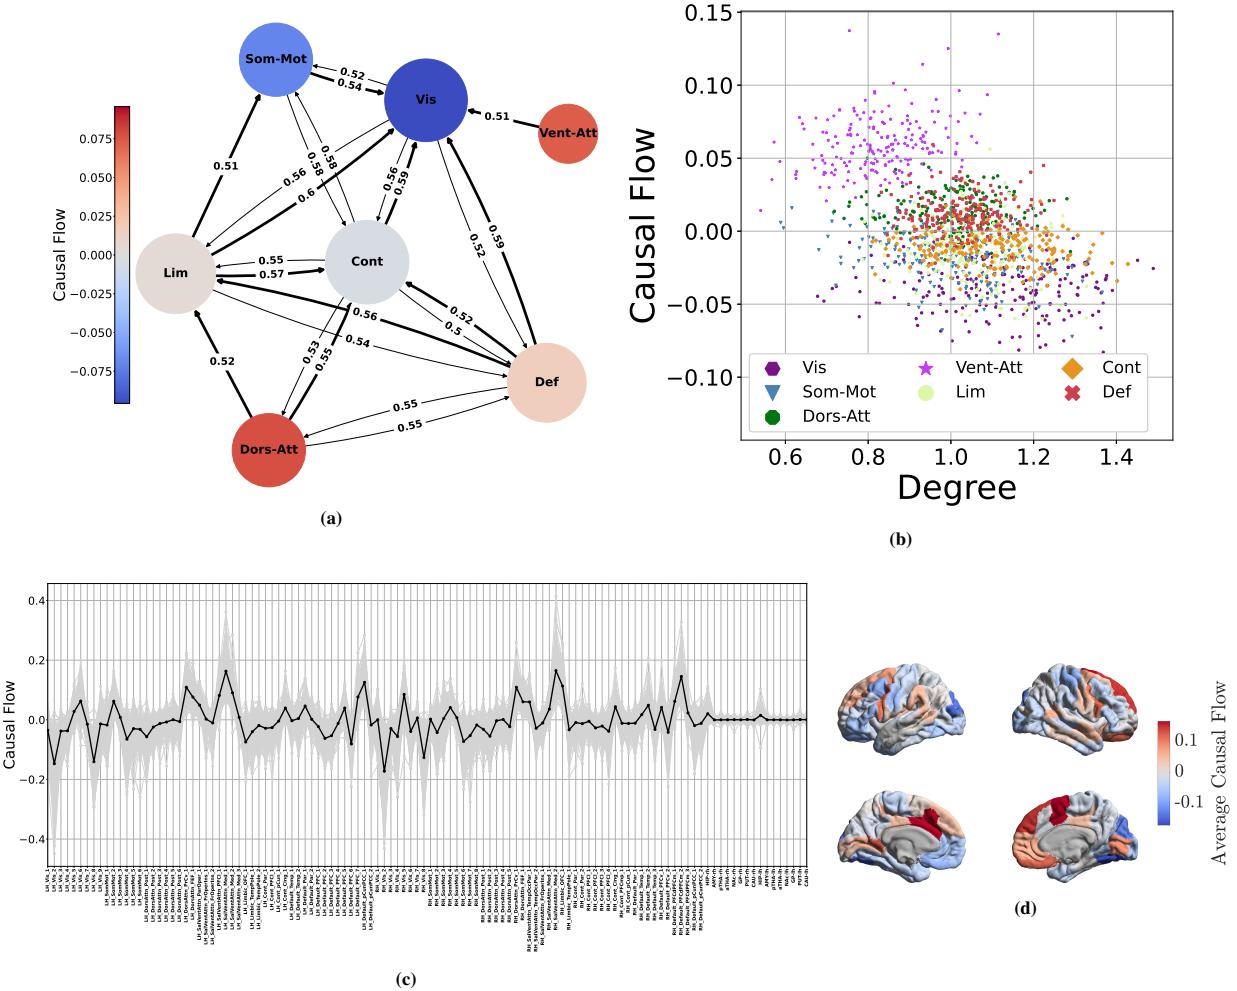


Figure 5: Patterns of causal flow across subjects. (a) The average subnetwork graph, computed as the mean of subnetwork graphs of all the subjects. In the subnetwork graph of each subject, the weight of an edge from subnetwork i to j is the number of nodes in subnetwork i that connect to nodes in subnetwork j , normalized by the number of all possible edges between these subnetworks. Edges with weights less than 0.5 are removed for better visualizations. To further ease the visual inspection of causal flows, if two networks are bidirectionally coupled we have shown the stronger edge with a thicker line (see Supplementary Figures 51, 52 for the weight matrix before thresholding and the corresponding nodal degrees and causal flows.) (b) The joint distributions of causal flow and degree for each "node" of the subnetwork graphs across all subjects. (c) Distribution of nodal causal flow, computed separately for each node in the causal graph and each subject. Gray colors correspond to different subjects and the black line shows the average nodal degrees across subjects. (d) Average cortical nodal causal flows (black line in (c)) shown over the brain cortical surface.

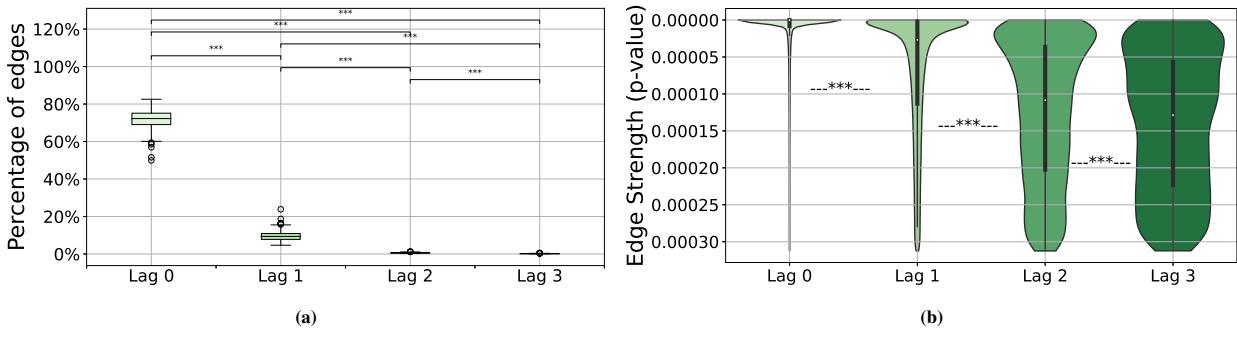


Figure 6: The contributions of each lag to the final causal graph in CaLLTiF. (a) For each lag, the box plot shows the distribution of the percentage of edges that come only from that lag across subjects. In all boxplots, the center line represents the median, the box spans the interquartile range (IQR), the whiskers extend up to 1.5 times the IQR from the box limits, and individual points beyond the whiskers indicate outliers. (b) The strength (statistical significance) of edges across lags. For each edge within the subgraph of each lag, we define its strength as the p-value of the partial correlation test that was used to conclude the presence of that edge (see Methods) even though all edges have a statistically significant p-value by definition, edges in larger lags are significantly closer to the threshold for significance than those in smaller lags. *** = $p < 0.001$, one-sided Wilcoxon rank-sum test.

315 lags had a non-zero (and significant by construction) contribution to the end graph in all subjects.
 316 Even lag 3 had a median of approximately 0.2% unique contributions to the final graph across
 317 subjects. We also found very small intersections among lags. This not only highlights the
 318 importance of considering multiple lags rather than just the first one or two but also demonstrates
 319 that it is incorrect to assume that if one region causes another, that causation will appear
 320 continuously across all lags. In summary, we found contemporaneous effects dominant in the
 321 final causal graphs of CaLLTiF, even though all lags had significantly non-zero and mostly unique
 322 contributions.

323 **Causal connections are modulated by pairwise Euclidean distance.** As one would expect
 324 from a network learned over a set of nodes embedded in physical space, the causal graphs learned
 325 by CaLLTiF are modulated in a number of ways by the Euclidean distance between pairs of
 326 nodes. First, we found degree similarity (correlation coefficient between nodal degrees of two
 327 parcels over all subjects) to decay statistically significantly, though weakly in effect size, with
 328 parcel distance (Pearson $r = -0.12$, $p = 10^{-43}$, 95% confidence interval $(-0.14, -0.1)$) as

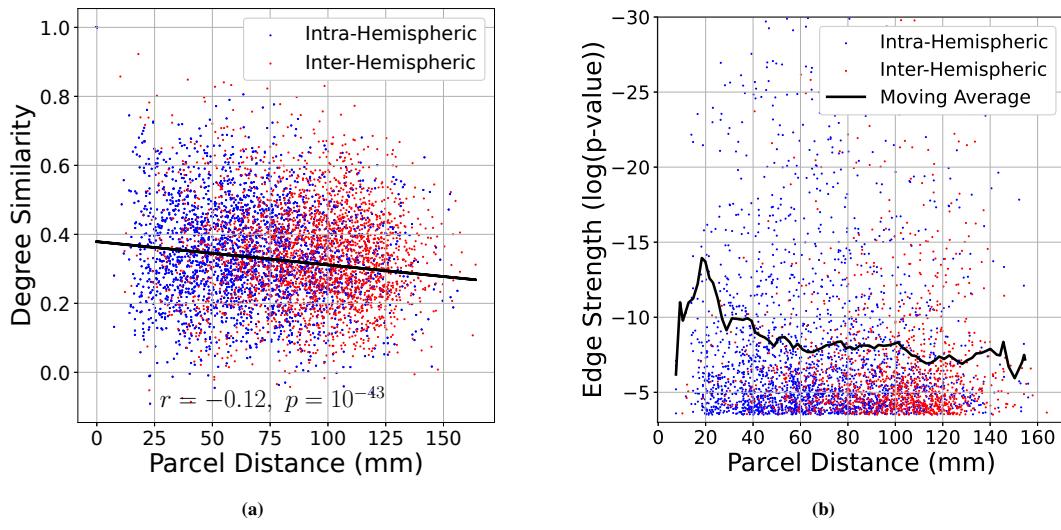


Figure 7: Effect of Euclidean distance on edge attributes. (a) Degree similarity (correlation coefficient between nodal degrees of two parcels over all subjects) as a function of the Euclidean distance between the parcels. Parcel pairs in the same hemisphere (intra-hemispheric) and parcels in two different hemispheres (inter-hemispheric) are shown in blue and red, respectively. Degree similarity decays statistically significantly with parcel distance (Pearson $r = -0.12$, $p = 10^{-43}$, 95% confidence interval $(-0.14, -0.1)$), much more so among intra-hemispheric parcels (Pearson $r = -0.27$, $p = 10^{-82}$, 95% confidence interval $(-0.29, -0.24)$) than inter-hemispheric ones (Pearson $r = -0.09$, $p = 10^{-5}$, 95% confidence interval $(-0.13, -0.05)$). (b) Edge strength (as in Figure 6b) as a function of Euclidean distance between edge endpoints (note the inverted vertical axis). The solid line shows the corresponding moving average of $\log(p)$ with 10mm window size and 8mm window overlap. The upper limit of the vertical axis is limited to -30 for better visualization.

329 shown in Figure 7a (See Supplementary Figures 54 and 55 for separate maps of degree
 330 similarities and pairwise nodal distances). This relationship is stronger among intra-hemispheric
 331 parcels (Pearson $r = -0.27$, $p = 10^{-82}$, 95% confidence interval $(-0.29, -0.24)$) where
 332 connections are denser and shorter-distance, compared to inter-hemispheric parcels (Pearson
 333 $r = -0.09$, $p = 10^{-5}$, 95% confidence interval $(-0.13, -0.05)$). Thus, in summary, nodes that
 334 are physically closer to each other also have more similar causal connections to the rest of the
 335 network, particularly if they belong to the same hemisphere.

336 The strength of CaLLTiF edges is also modulated by the Euclidean distance between edge
 337 endpoints, even though we observed that there are approximately as many long-distance edges as

338 short ones (See Supplementary Figures 56). We define the strength of each edge in the final graph
339 (union over lags) as the *minimum* p-value of respective partial correlation tests across all lags (cf.
340 Methods). As seen from Figure 7b, the mean strength of causal edges (black solid line) initially
341 increases with the Euclidean length of the edge until about 20mm and then decays with Euclidean
342 edge length thereafter.

343 Finally, we found no major differences between the Euclidean distances of edges contributed
344 by different lags. Given that causal effects take time to spread along axonal fibers throughout the
345 brain, one might expect physically-closer pairs of nodes to be connected by lower-lag edges and
346 more distant pairs of nodes to be connected by larger-lag edges. However, as seen in
347 Supplementary Figure 58, this is not quite the case. Given the slow sampling of fMRI, even the
348 most distant regions can causally affect each other in time scales shorter than 1 TR. Thus, the
349 observation that the physical distance of pairs of nodes was not related to edge lag should not be
350 taken as an indication that such relationships would – or would not – be observed when sampling
351 with higher temporal precision.

352 **Degree, but not causal flow, shows significant laterality and gender differences.** We
353 observed that nodal degrees were statistically significantly higher in the right hemisphere
354 (Figure 8a, Cohen's $d = 0.07$ and $p = 10^{-48}$, one-sided Wilcoxon signed-rank test), even though
355 no such laterality was found in nodal causal flows (Figure 8b, Cohen's $d = 0.02$ and $p = 0.23$,
356 one-sided Wilcoxon signed-rank test). To understand which subnetworks might be playing a
357 stronger role in the hemispheric asymmetry observed in the distribution of nodal degrees,
358 Figure 8c shows the mean degrees of corresponding pairs of regions in the left and right
359 hemispheres, color-coded by functional subnetworks (cf. Supplementary Figures 49 for separate
360 plots per subnetwork). The ventral attention, dorsal attention, and executive control networks
361 show clearly larger causal degrees in the right hemisphere, whereas the limbic network and DMN
362 have larger causal degrees in the left hemisphere. A similar plot for causal flows (Figure 8d,
363 Supplementary Figure 50) shows a lot more symmetry, except for the limbic network which
364 shows exceptionally higher causal flows (i.e., source-ness) in the right compared to the left
365 hemispheres. The DMN also shows some asymmetry in its causal flow, where right DMN nodes
366 are mostly sources of causal flow whereas left DMN causal flows are more evenly distributed

367 around zero. Thus, in summary, various functional subnetworks show laterality in degree
368 distributions, culminating in a net increase in right vs. left nodal degrees. Causal flows, however,
369 are mostly symmetric, except for the limbic network which shows a strong flow from the right to
370 the left hemisphere.

371 Similarly, degree, but not causal flow, shows a small but statistically significant difference
372 between men and women. In causal connectomes learned by CaLLTiF, we found nodal degrees to
373 be statistically significantly higher in women compared to men (Figure 8e, Cohen's $d > 0.05$,
374 $p < 10^{-5}$, one-sided Wilcoxon rank-sum test). Nodal causal flows, on the other hand, were
375 statistically indistinguishable between men and women (Figure 8f, Cohen's $d = -3.77 \times 10^{-18}$,
376 $p = 0.81$, one-sided Wilcoxon rank-sum test). These result demonstrate that nodal degrees in
377 causal connectomes are generally more heterogeneous and sensitive while causal flows tend to be
378 more homogeneous and stereotyped across individuals and hemispheres. Further research is
379 needed to pinpoint the root causes of these differences (and lacks thereof) as well as potential
380 implications of them in health and disease.

DISCUSSION

381 In this study, we investigated the problem of whole-brain causal discovery from fMRI. We first
382 comprehensively compared existing causal discovery techniques suitable for whole-brain fMRI by
383 examining both theoretical properties and numerical outcomes on simulated fMRI. To address the
384 limitations of existing algorithms, we proposed CaLLTiF which improves upon the state of the art
385 in several directions including learning contemporaneous edges and cycles, type I error control,
386 and scalability. A core aspect of CaLLTiF is its treatment of contemporaneous effects. Our results
387 with the HCP data (Figures 6a and 6b) confirmed the importance of being able to reveal such
388 “contemporaneous” effects, where these effects accounted for the majority and strongest of
389 network edges. Further, the distributions of edges with different Euclidean distances at each lag
390 (Supplementary Figure 58) demonstrates how broadly neural signals can propagate across the
391 brain in one TR interval, even with the relatively fast sampling (TR = 0.72s) in the HCP dataset.

392 Furthermore, in interpreting CaLLTiF's outputs, it is important to note its by-design
393 conservative method of correction for temporal multiple comparisons. In the Macaque_Full

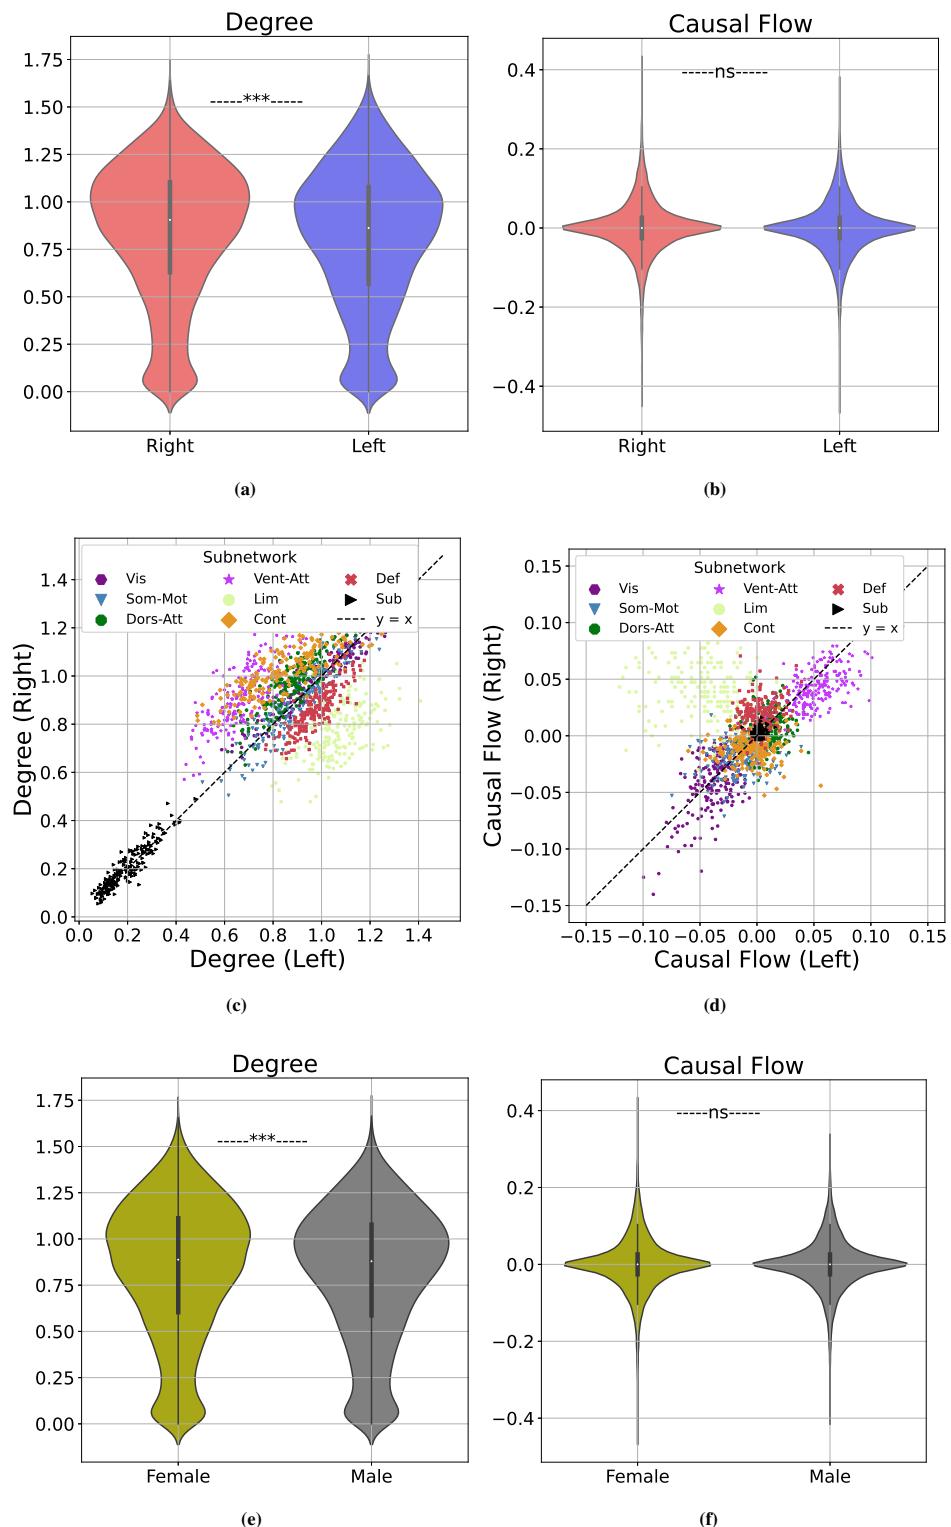


Figure 8: Hemispheric laterality and gender differences in causal connectomes. (a) Distributions of nodal degree in the right and left hemispheres, combined across all subjects. (b) Similar to (a) but for causal flows. (c) Nodal degrees, averaged across subjects and color-coded by functional subnetwork, for pairs of corresponding parcels within the right and left hemispheres. To properly pair nodes across the two hemispheres, degrees of all the parcels with the same label in the Schaefer/Tian atlas were averaged and then paired. Deviations from the dashed $y = x$ line indicate laterality. (d) Similar to (c) but for causal flows. (e) Distributions of nodal degrees in female and male subjects, combined across all brain regions. (f) Similar to (e), but for causal flows.

394 simulated dataset where the ground truth is known, we found Alpha Level = 0.01 to maximize the
395 F1 score, while CaLLTiF's correction for temporal multiple comparisons would have suggested
396 $0.01/32 = 0.0003$ (cf. Equation (4)) and thus would have obtained sparser graphs. Similarly, we
397 obtained causal graphs from human fMRI that are about 30-55% dense across all subjects and
398 40-60% dense among cortical nodes (Supplementary Figure 43a). These are generally consistent
399 with, but sparser than, the near-66% cortical density found using tract-tracing results in
400 non-human primates (Markov et al., 2014). In other words, graphs returned by CaLLTiF are likely
401 to have higher precision but lower recall than what would maximize the F1 score. This
402 conservativeness is by design and desirable *in the absence of a ground-truth causal graph*,
403 ensuring that we have strong confidence in all discovered edges (at least 99% confidence in every
404 detected edge in our experiments with human fMRI). Nevertheless, CaLLTiF's level of
405 conservativeness can also be tuned as needed by tuning its pre-correction significance threshold (q
406 in Equation (4)).

407 An unexpected finding of our study was the higher accuracy of causal discovery when
408 conditioning pairwise independence tests (see Equation (1)) on all other nodes in the network, as
409 done in CaLLTiF, compared to using a more restricted parent set found by PCMCI (cf.
410 Supplementary Figure 16). The approach taken by PCMCI increases statistical power (cf. the
411 trend of optimal 'Alpha Level' values in Supplementary Figure 16b), but can significantly
412 increase type I error in the presence of contemporaneous effects. Even further, we found that even
413 using the (lagged) ground-truth parent sets for each node leads to a lower F1 score compared to
414 using complete conditioning sets (Supplementary Figure 39-41). This is likely because
415 CaLLTiF's conditioning on the *past* of all variables serves as a proxy for the missing
416 contemporaneous parents that should have been conditioned on. On the other hand, one may
417 wonder if this issue could have been better resolved by conditioning on contemporaneous
418 variables themselves. However, this can result in spurious statistical dependence if conditioning
419 on all contemporaneous variables (consider, e.g., testing $X_i \not\perp\!\!\!\perp X_j | X_k$ with the ground-truth
420 causal graph $X_i \rightarrow X_k \leftarrow X_j$). For a detailed discussion on this see Supplementary Note 3.

421 **Causal connectivity during resting state.** When applying CaLLTiF to resting state human
422 fMRI, we found the strongest causal effect to flow from attention to sensorimotor networks. The

423 strongest sources and sinks were the ventral attention and visual networks, followed by dorsal
424 attention and somatomotor networks, respectively. Despite the lack of “ground-truth”
425 connectivity as in simulated data, we can still evaluate these findings based on their agreement
426 with prior findings on the roles of resting state networks. The dorsal attention network (involving
427 regions in the intraparietal and superior frontal cortex) is commonly believed to handle top-down
428 selection processes and is only modulated by stimulus detection, while the ventral attention
429 network (including areas such as the temporoparietal and inferior frontal cortex) specializes in
430 detecting behaviorally relevant stimuli, particularly those that are salient or unexpected, and
431 directing attention to them (Corbetta & Shulman, 2002). These networks exhibit internally
432 correlated activity patterns (particularly during the task) (Corbetta, Patel, & Shulman, 2008) and
433 their flexible interaction facilitates dynamic attentional control aligned with both top-down goals
434 and bottom-up sensory inputs (Vossel, Geng, & Fink, 2014). Nonetheless, the ventral attention
435 network is more strongly involved in the “reorientation” of attention, namely, interrupting one
436 thought process and orienting attention towards something salient, while the dorsal attention
437 network is more strongly implicated in focused and guided attentional tasks such as visual search
438 under high short-term memory load (Vossel et al., 2014).

439 Arguably, resting state activity is more aligned with the former (salience-based reorientation)
440 than the latter (goal-driven focused attention). Despite a lack of sensory salience, attention is
441 frequently reoriented during periods of rest by endogenously-salient thoughts and memories.
442 Intermittent periods of focused attention can also arise, particularly given the long durations of
443 each resting-state scan (~ 15min). Our findings thus confirm and corroborate the existing
444 hypotheses about the roles of attention networks and how they jointly but unevenly drive brain
445 dynamics during rest. Furthermore, due to the lack of meaningful sensory (particularly visual)
446 input during a resting state scan, sensory areas receive more top-down influence from
447 higher-order cortices than they provide bottom-up information to them. As such, the sink-ness of
448 sensorimotor areas in graphs learned by CaLLTiF is arguably more consistent with the nature of
449 resting state activity compared to a contrasting, sensory-driven flow found, e.g., in (Rawls et al.,
450 2022). Finally, we found the DMN to be both a hub and a weak source of causal flow, which is

451 consistent with its well-known role in resting state dynamics (Andrews-Hanna, Smallwood, &
452 Spreng, 2014; Greicius, Krasnow, Reiss, & Menon, 2003; Raichle et al., 2001).

453 Resting-state causal graphs learned by CaLLTiF are also notably aligned with the literature in
454 terms of the laterality of different functional subnetworks (Figure 8c and Supplementary
455 Figure 49). Several studies have found the ventral attention system to be predominantly
456 lateralized to the right hemisphere (Corbetta et al., 2008; Corbetta & Shulman, 2002; Mengotti,
457 Käsbauer, Fink, & Vossel, 2020; Vossel et al., 2014). Similarly, the degree distribution of ventral
458 attention nodes in graphs learned by CaLLTiF is strongly right-lateralized. We found the dorsal
459 attention network to also be right-lateralized, but not as strongly as the ventral network. Similarly,
460 The dorsal attention network is found by prior studies to be organized mostly bilaterally, except
461 for specific regions (Intraparietal sulcus and frontal eye field) in the right dorsal network that
462 show stronger involvement in the attentional control of both hemispheres compared to their left
463 counterparts (Mengotti et al., 2020). Also similar is alignment in the lateralization of the default
464 mode network, where both CaLLTiF and several past studies have found it to be left
465 lateralized (Agcaoglu, Miller, Mayer, Hugdahl, & Calhoun, 2015; Banks et al., 2018; Nielsen,
466 Zielinski, Ferguson, Lainhart, & Anderson, 2013; Swanson et al., 2011). Sensorimotor cortices,
467 on the other hand, were found to be highly symmetric and not particularly lateralized to either
468 hemisphere in causal graphs learned by CaLLTiF, a finding that is also consistent with the
469 generally symmetric involvement of primary sensory and motor areas in contralateral
470 processing (Agcaoglu et al., 2015). Finally, research on the laterality of the control and limbic
471 networks is ongoing and, to the best of our knowledge, inconclusive (see, e.g., (Morton, 2020)).
472 In graphs learned by CaLLTiF, however, we observe strong left lateralization of the limbic and
473 right lateralization of the control networks, respectively. Thus, in summary, we observe clear
474 laterality in all but sensorimotor cortical networks which either corroborate the existing literature
475 or clarify previously inconclusive observations.

476 **Hyperparameter selection and sensitivity.** CaLLTiF has two hyperparameters, τ_{\max} and
477 α_{level} . Due to CaLLTiF's type-I error control over lags (see Methods), these parameters are
478 interdependent, as described by Equation (4). This makes τ_{\max} the only effectively 'free'
479 hyper-parameter, which can be systematically selected by analyzing how much each additional

lag contributes to the final graph (as in Figure 6a). Furthermore, the end results of CaLLTiF are remarkably robust to variations in τ_{\max} . Supplementary Figure 59 illustrates the sensitivity of CaLLTiF by comparing graphs generated from real HCP data with $\tau_{\max} = 1, 2$, and 4 against those with $\tau_{\max} = 3$ analyzed in the main text. For comparison, we also show the percentage of changes in resulting graphs if we fix the value of α_{level} (i.e., ignore Equation (4)) so that, e.g., increasing τ_{\max} from 3 to 4 only adds lag-4 edges without affecting other lags. As we can see, at $\tau_{\max} = 2$ and 4 (33% change in τ_{\max}), the resulting graphs change less than 6% for fixed type I error and less than 3% for fixed α_{level} . Similarly, when changing τ_{\max} to 1 (67% change in τ_{\max}), the resulting graphs change by less than 13% and 6% in the two conditions, respectively.

Limitations. The present study has a number of limitations. From a biological perspective, synaptic causation happens at the level of neuronal activity, from which fMRI is a noisy readout. This lack of access to the true underlying neural activity limits the accuracy of any causal discovery method working with fMRI, and attempting to extract the underlying neural activity from fMRI data is often futile (Supplementary Figure 42, also cf. (Nozari et al., 2023)). The low temporal resolution of fMRI, even with the TR value of 720ms in the HCP data, also limits the precision of causal discovery. As we saw from Supplementary Figure 58, edges of all lengths are observed even at lag 0. This indicates the possibility that some of the edges discovered by CaLLTiF may be polysynaptic paths but resemble a direct monosynaptic connection at low temporal resolution. Finally, similar to most constraint-based methods, the causal graphs returned by CaLLTiF are not tied to a generative dynamical model (as is the case with VARLiNGAM, DYNOTEARs, DCM, etc). If such generative models are needed, VAR models based on CaLLTiF's extended time-lagged graph constitute a natural choice, but further research is needed to compare the dynamic predictive accuracy of such models against potential alternatives (Ljung, 1999).

Conclusions. Overall, this study demonstrates the interplay between the theoretical challenges of causal discovery and the practical limitations of fMRI, and the design of an algorithmic solution that can bridge this gap. This work motivates several follow-up studies, including the application of the proposed CaLLTiF method to task fMRI and comparing its outcomes against structural connectivity.

MATERIAL AND METHODS

509 *Simulated fMRI Data*

510 When comparing different causal discovery algorithms or different hyperparameters of the same
511 algorithm, we used several benchmarks of simulated fMRI data with known ground truth
512 connectivity from (Sanchez-Romero et al., 2019). In general, this dataset included two groups of
513 networks, one consisting of 9 simple small-scale synthetic graphs and one consisting of two
514 graphs extracted from the macaque connectome. From the latter group, we only used the smallest
515 (Macaque_SmallDegree) and the largest (Macaque_Full).

516 The details of generating BOLD signals from each graph are detailed in (Sanchez-Romero et
517 al., 2019). In brief, the same simulation procedure was used for simple and macaque-based
518 graphs, where the authors used the model proposed in (Smith et al., 2011) which is itself based on
519 the DCM architecture of (K. Friston et al., 2003). Underlying neural dynamics are simulated
520 using the linear differential equation $dz/dt = \sigma Az + Cu$, where A denotes the ground-truth
521 connectivity. To simulate resting-state data, the u input was modeled using a Poisson process for
522 each of the regions ($C = I$). The neuronal signals z were then passed through the
523 Balloon-Windkessel model (Buxton, Wong, & Frank, 1998; Smith et al., 2011) to obtain
524 simulated BOLD data.

525 *Resting-State fMRI from the Human Connectome Project*

526 For the real fMRI analysis, we used ICA-FIX resting-state data from the Human Connectome
527 Project S1200 release (Barch, 2017; Burgess et al., 2016; Essen et al., 2013). Resting-state fMRI
528 images were collected with the following parameters: TR = 720 ms, TE = 33.1 ms, flip angle = 52
529 deg, FOV = 208x108 mm, matrix = 104x90, slice thickness = 2.0 mm, number of slices = 72 (2.0
530 mm isotropic), multi-factor band = 8, and echo spacing = 0.58 ms. Brains were normalized to
531 fslr32k via the MSM-AII registration and the global signal was removed. We removed subjects
532 from further analysis if any of their four resting state scans had excessively large head motion,
533 defined by having frames greater than 0.2 mm frame-wise displacement or a derivative root mean
534 square (DVARS) above 75. Also, subjects listed in (Elam, 2020) under “3T Functional
535 Preprocessing Error of all 3T RL fMRI runs in 25 Subjects” or “Subjects without Field Maps for

536 “Structural scans” were removed. Among the remaining 700 subjects, the 200 with the smallest
537 head motion (DVARS) were selected for analysis. For all subjects, we parcellated the brain into
538 100 cortical regions (Schaefer 100x7 atlas (Schaefer et al., 2018)) and 16 subcortical ones
539 (Melbourne Scale I atlas (Tian, Margulies, Breakspear, & Zalesky, 2020)). The Human
540 Connectome Project experiments were carried out by the WU-Minn consortium and its adherence
541 to ethical standards was approved by the by the Internal Review Board of the respective
542 institutions. Explicit informed consent was acquired from all participants involved in the
543 study (Essen et al., 2013).

544 ***Causal discovery methods***

545 One aim of causal inference is to construct a causal graph based on observational data. The
546 relationship between a probability distribution and its depiction as a graph plays a significant role
547 in this process. Nevertheless, it is not always feasible to deduce a causal graph solely from
548 observational data. Further assumptions are therefore required. Here, we briefly summarize the
549 main assumptions and principles underlying the list of causal discovery methods studied in this
550 work (cf. Table 1).

551 **PCMCI.** PCMCI was proposed in (Runge et al., 2019) as a constraint-based causal discovery
552 method designed to work with time-series data. The algorithm is composed of two main steps. In
553 the first step, the algorithm selects relevant variables using a variant of the undirected graph
554 discovery part of the PC algorithm (Spirtes & Glymour, 1991). This step removes irrelevant
555 variables for conditioning and therefore increases statistical power. In the second step, the
556 algorithm uses the momentary conditional independence (MCI) test, which measures the
557 independence of two variables conditioned on the set of their parents identified in step 1. The
558 MCI test helps to reduce the false positive rate, even when the data is highly correlated. PCMCI
559 assumes that the data is stationary, has time-lagged dependencies, and has causal sufficiency.
560 Even when the stationarity assumption is violated, PCMCI was shown to perform better than
561 Lasso regression or the PC algorithm (Runge et al., 2019). However, PCMCI is considered not
562 suitable for highly predictable (almost deterministic) systems with little new information at each

563 time step (Runge et al., 2019). The Python implementation of PCMCI is available in the

564 Tigramite package at <https://github.com/jakobrunge/tigramite>.

565 As noted earlier, PCMCI only returns $\circ-\circ$ edges among contemporaneous variables. While this

566 allows PCMCI to relax the common DAG assumption and allow for cycles, it results in a mixed

567 summary graph, where multiple types of edges (\leftarrow , \rightarrow , and/or $\circ-\circ$) can exist between two nodes.

568 In contrast, we require all algorithms to output a directed graph. Therefore, when reporting F1

569 scores for PCMCI, we only include directed edges coming from lagged relationships and exclude

570 the contemporaneous $\circ-\circ$ edges. The only exception is what we call ‘Mixed PCMCI’ (See

571 Supplementary Figures 28- 30), where the contemporaneous $\circ-\circ$ edges are also included in the

572 computation of *adjacency* F1 scores.

573 **PCMCI⁺.** PCMCI⁺ is an extension of the PCMCI method which incorporates directed

574 contemporaneous links in addition to the lagged ones (Runge, 2020). The approach revolves

575 around two key concepts. First, it divides the undirected graph edge removal phase into separate

576 lagged and contemporaneous conditioning phases, thereby reducing the number of conditional

577 independence tests required. Second, it incorporates the idea of momentary conditional

578 independence (MCI) tests from PCMCI (Runge et al., 2019) specifically in the contemporaneous

579 conditioning phase. PCMCI⁺ also outputs a time-series graph with different types of

580 contemporaneous edges, including directed edges (\rightarrow and \leftarrow), unoriented edges ($\circ-\circ$), and

581 conflicting edges ($\times - \times$). Consistent with our requirement of a regular digraph at the end, we

582 disregarded the unoriented and conflicting edges and retained only the directed ones. Similar to

583 most other causal discovery algorithms, PCMCI⁺ does not permit cycles in the contemporaneous

584 links, which could potentially account for its relatively underwhelming performance over fMRI

585 data. The Python implementation of PCMCI+ is also available in the Tigramite package

586 <https://github.com/jakobrunge/tigramite>.

587 **VARLiNGAM.** VARLiNGAM is a causal discovery method that combines non-Gaussian

588 instantaneous models with autoregressive models. This method, proposed in (Hyvärinen et al.,

589 2010), builds on the fact that in the absence of unobserved confounders, linear non-Gaussian

590 models can be identified without prior knowledge of the network structure. VARLiNGAM is

591 capable of estimating both contemporaneous and lagged causal effects in models that belong to

592 the class of structural vector autoregressive (SVAR) models and provides ways to assess the
593 significance of the estimated causal relations. These models are a combination of structural
594 equation models (SEM) and vector autoregressive (VAR) models. In addition, VARLiNGAM
595 emphasizes the importance of considering contemporaneous influences, as neglecting them can
596 lead to misleading interpretations of causal effects. Nevertheless, VARLiNGAM does not permit
597 cycles in the contemporaneous links either, which could potentially account for its relatively poor
598 performance over brain fMRI data with many feedback loops. The VARLiNGAM method is
599 available from <https://github.com/cdt15/lingam> and a tutorial can be found at
600 <https://lingam.readthedocs.io/en/latest/tutorial/var.html>.

601 **DYNOTEARS.** Dynamic NOTEARS (DYNOTEARS) method, proposed in (Pamfil et al.,
602 2020), is a score-based method designed to discover causal relationships in dynamic data. It
603 simultaneously estimates relationships between variables within a time slice and across different
604 time slices by minimizing a penalized loss function while ensuring that the resulting directed
605 graph is acyclic (including acyclicity of contemporaneous connections). The goal is to identify
606 the best set of conditional dependencies that are consistent with the observed data. DYNOTEARS
607 builds on the original NOTEARS method proposed in (Zheng, Aragam, Ravikumar, & Xing,
608 2018), which uses algebraic properties to characterize acyclicity in directed graphs for static data.
609 Python implementations are available from the CausalNex library
610 (<https://github.com/quantumblacklabs/causalnex>) as well as
611 https://github.com/ckassaad/causal_discovery_for_time_series.

612 **DGlearn.** DGlearn is a score-based method for discovering causal relationships from
613 observational data. Importantly, it is one of few algorithms that can learn cyclic structures from
614 cross-sectional data. The method, introduced in (Ghassami et al., 2020), is based on a novel
615 characterization of equivalence for potentially cyclic linear Gaussian directed graphical models.
616 Two structures are considered equivalent if they can generate the same set of data distributions.
617 DGlearn utilizes a greedy graph modification algorithm to return a graph within the equivalence
618 class of the original data-generating structure. The Python implementation of DGlearn is
619 available at <https://github.com/syanga/dglearn>.

620 **FASK.** The Fast Adjacency Skewness (FASK) method, proposed in (Sanchez-Romero et al.,
621 2019), is a hybrid method for causal discovery from cross-sectional data, combining
622 constraint-based and noise-based elements. It leverages (and needs) non-Gaussianity in the data
623 and allows for cycles. This algorithm is composed of two main steps. The first step, called
624 FAS-Stable, outputs an undirected graph G_0 by iteratively performing conditional independence
625 tests under the increasing size of the conditioning set and using the Bayesian information criterion
626 (BIC) to compare the conditioning sets. In the second step, assuming i.i.d. non-Gaussian data,
627 each of the $X - Y$ adjacencies in G_0 are oriented as a 2-cycle (\leftrightarrow) if the difference between
628 $\text{corr}(X, Y)$ and $\text{corr}(X, Y|X > 0)$, and $\text{corr}(X, Y)$ and $\text{corr}(X, Y|Y > 0)$, are both
629 significantly nonzero, and as a unidirectional edge otherwise. The pseudo-code for FASK can be
630 found in Supporting Information A of (Sanchez-Romero et al., 2019) and Java source code for it
631 is available at <http://github.com/cmu-phil/tetrad>.

632 **MVGC.** In (Granger, 1969), Granger introduced a statistical version of Hume's regularity
633 theory, stating that X_p Granger-causes X_q , if past values of X_p provide unique, statistically
634 significant information about future values of X_q (Assaad et al., 2022). While this allows for
635 optimal forecasting of an effect and has been extended to multivariate systems (Barnett & Seth,
636 2014), MVGC cannot account for contemporaneous effects and the presence of unobserved
637 confounders can result in spurious edges. Python implementation of MVGC is available at
638 https://github.com/ckassaad/causal_discovery_for_time_series.

639 **NTS-NOTEARS.** NTS-NOTEARS is a nonlinear causal discovery method designed for
640 time-series data (Sun, Liu, Poupart, & Schulte, 2021). It employs 1-D convolutional neural
641 networks to capture various types of relationships, including linear, nonlinear, lagged, and
642 contemporaneous connections among variables. The method ensures that the resulting causal
643 structure forms a directed acyclic graph. It builds upon the NOTEARS approach (Zheng et al.,
644 2018), and is similarly based on continuous optimization. Similar to other algorithms above, it
645 assumes the presence of no hidden confounding factors and stationarity of the data-generating
646 process. In our analysis, we compare NTS-NOTEARS as a state-of-the-art nonlinear method
647 against the aforementioned linear algorithms in synthetic fMRI (cf. Supplementary Figure 38). A

648 Python implementation of NTS-NOTEARS is available at

649 <https://github.com/xiangyu-sun-789/NTS-NOTEARS>

650 **CaLLTiF (proposed method).** The proposed CaLLTiF method builds upon PCMCI (Runge et
651 al., 2019) but, instead of using a PC-type approach in the first step to estimate the set of parents
652 for lagged variables, it starts from a complete conditioning set including all lagged variables. This
653 choice dramatically decreases computational cost, but surprisingly, it is also optimal, as shown in
654 Supplementary Figure 16, because as mentioned in the discussion section, the approach of
655 PCMCI discards contemporaneous effects. Using a complete conditioning set, CaLLTiF then
656 performs Momentary Conditional Independence (MCI) partial correlation tests between all pairs
657 of variables. Specifically, for any pair $X_i(t - \tau), X_j(t)$ with $i, j \in 1, \dots, N$ and time delays
658 $\tau \in 0, 1, \dots, \tau_{\max}$, a causal link is established $(X_i(t - \tau) \rightarrow X_j(t))$ if $\tau > 0$ and $X_i(t) \circ\circ X_j(t)$ if
659 $\tau = 0$), if and only if:

$$\text{Cond. Ind. Test : } X_i(t - \tau) \not\perp\!\!\!\perp X_j(t) \mid \{X_k(t - s) : k = 1, \dots, N, s = 1, \dots, \tau_{\max}\} \setminus \{X_i(t - \tau)\}. \quad (1)$$

660 Note that, despite being complete, the conditioning sets only include variables from *prior* time
661 lags. As noted earlier, to test a conditional independence of the form $X \not\perp\!\!\!\perp Y | Z$, we compute the
662 partial correlation coefficient $\rho(X, Y | Z)$ between X and Y conditioned on the set of variables in
663 Z and the corresponding p-value for the null hypothesis that $\rho(X, Y | Z) = 0$. An edge is placed
664 between $X_j(t)$ and $X_i(t - \tau)$ if this p-value is less than the hyperparameter ‘Alpha Level’. The
665 value of this threshold was selected optimally in simulated fMRI and using temporal correction
666 for multiple comparisons (see below) in real data. Finally, for contemporaneous pairs ($\tau = 0$),
667 each $\circ\circ$ edge is replaced with \leftrightharpoons if there are no other edges between those two variables at other
668 lags, and is replaced with a directed edge or a \leftrightharpoons based on the lagged direction(s) otherwise. For a
669 more detailed summary of the partial correlation-based edge discovery in CaLLTiF, see
670 Supplementary Note 2. A pseudocode of CaLLTiF is shown in Algorithm 1.

Algorithm 1 Causal discovery for Large-scale Low-resolution Time-series with Feedback (CaLLTiF)

Input: $X_i(t)$ for $i = 1, \dots, N$ and $t = 1, \dots, T$

Output: Summary causal graph

Hyperparameters: τ_{\max} : Maximum Time Lag, α_{level} : Per-lag threshold for statistical significance

```
1: Initialize the extended causal graph with nodes  $\{X_i(t - \tau) \mid i = 1, \dots, N, \tau = 0, \dots, \tau_{\max}\}$ 
   and no edges
2: Initialize the summary causal graph with nodes  $\{X_i \mid i = 1, \dots, N\}$  and no edges
3: Build a complete conditioning set:  $S_{\text{cond}} \leftarrow \{X_k(t - s) \mid k = 1, \dots, N, s = 1, \dots, \tau_{\max}\}$ 
4: for  $i \in \{1, \dots, N\}$  do
5:   for  $j \in \{1, \dots, N\}$  do
6:     for  $\tau \in \{0, 1, \dots, \tau_{\max}\}$  do
7:       Run Conditional Independence Test:  $X_i(t - \tau) \perp X_j(t) \mid S_{\text{cond}} \setminus \{X_i(t - \tau)\}$ 
8:       if p-value  $< \alpha_{\text{level}}$  then
9:         if  $\tau > 0$  then
10:          Add edge  $X_i(t - \tau) \rightarrow X_j(t)$  to the extended causal graph
11:        else
12:          Add edge  $X_i(t) \circ\circ X_j(t)$  to the extended causal graph
13:        end if
14:      end if
15:    end for
16:    if At least one lagged edge exists from  $X_i$  to  $X_j$  then
17:      Place a directed edge  $X_i \rightarrow X_j$  in the summary causal graph
18:    else if The contemporaneous edge  $X_i(t) \circ\circ X_j(t)$  exists then
19:      Place a two-cycle  $X_i \leftrightharpoons X_j$  in the summary causal graph
20:    end if
21:  end for
22: end for
```

671 Finally, it is imperative to acknowledge the possibility that some of the directed edges detected
672 by our methodology do not possess a strictly causal connotation. As previously indicated, the
673 orientation method relies on the widely accepted premise that bidirectional connections hold
674 notably greater prevalence than unidirectional links. Thus, we believe that the presented approach
675 shall yield a proximate representation of the true causal graph, concurrently accommodating
676 cyclic structures and circumventing computational barriers.

677 ***Construction of summary causal graphs from causal graphs over lagged variables***

678 Causal discovery algorithms designed for time series data often return a causal graph among the
679 lagged variables

$$\begin{aligned} & X_1(t - \tau_{\max}), \dots, X_1(t - 1), X_1(t) \\ & X_2(t - \tau_{\max}), \dots, X_2(t - 1), X_2(t) \\ & \vdots \\ & X_n(t - \tau_{\max}), \dots, X_n(t - 1), X_n(t) \end{aligned} \tag{2}$$

680 For algorithms other than CaLLTiF, from this we extract a final *summary* graph among the
681 variables X_1, \dots, X_n by placing an edge from X_i to X_j if there exists any $\tau \geq 0$ for which there
682 is an edge from $X_i(t - \tau)$ to $X_j(t)$. This is equivalent to an OR operation among binary edges (as
683 opposed, e.g., to a majority vote) and must be taken into account when interpreting the obtained
684 summary graphs. The process is similar in CaLLTiF except that the direction(s) of
685 contemporaneous $\circ-\circ$ edges are first resolved using lagged edges before executing the OR across
686 lags (cf. Algorithm 1).

687 ***Correction for multiple comparisons across lags in CaLLTiF***

688 As noted above, CaLLTiF places an edge from X_i to X_j in its summary graph if there exists at
689 least one $\tau \geq 0$ for which there is an edge from $X_i(t - \tau)$ to $X_j(t)$. Therefore, the decision to
690 place an edge from X_i to X_j depends on the outcomes of $\tau_{\max} + 1$ statistical tests, and to maintain
691 a desired bound on the probability of type I error for each edge in the *summary* graph, we need to
692 account for multiple comparisons across lags.

693 Formally, for each edge $X_i \rightarrow X_j$ in the final graph, the null hypothesis (i.e., lack of a direct
 694 causal effect from X_i to X_j) can be formulated as

$$H_0 = H_{0,0} \cap H_{0,1} \cap \cdots \cap H_{0,\tau_{\max}} = \bigcap_{\tau=0}^{\tau_{\max}} H_{0,\tau},$$

$$H_{0,\tau} = \{X_i(t - \tau) \not\rightarrow X_j(t)\}.$$

695 Let p_τ denote the p-value of the partial correlation test between $X_i(t - \tau)$ and $X_j(t)$ and α_{level}
 696 denote the significance threshold for each partial correlation test. Then, the probability of type I
 697 error is

$$P(\text{Type I Error}) = P\left(\bigcup_{\tau=0}^{\tau_{\max}} \{p_\tau < \alpha_{\text{level}}\} \mid \bigcap_{\sigma=0}^{\tau_{\max}} H_{0,\sigma}\right)$$

698 Note that this is different from the family-wise error rate (FWER, bounded by the Bonferroni
 699 method and its extensions) or the false discovery rate (FDR). In particular, this is different from
 700 FWER in that only one decision is made and the probability of error is computed for that single
 701 decision only. So, for instance, if in reality any subset (even one) of $\{H_{0,\tau}\}$ is false and the
 702 algorithm rejects any subset (even all) of $\{H_{0,\tau}\}$, there is no type I error, as an edge exists from
 703 X_i to X_j both in the data-generating process and in the final summary graph.

704 The type I error can then be bounded as

$$P(\text{Type I Error}) = P\left(\bigcup_{\tau=0}^{\tau_{\max}} \{p_\tau < \alpha_{\text{level}}\} \mid \bigcap_{\sigma=0}^{\tau_{\max}} H_{0,\sigma}\right) \leq \sum_{\tau=0}^{\tau_{\max}} P\left(\{p_\tau < \alpha_{\text{level}}\} \mid \bigcap_{\sigma=0}^{\tau_{\max}} H_{0,\sigma}\right)$$

$$= \sum_{\tau=0}^{\tau_{\max}} \frac{P\left(\{p_\tau < \alpha_{\text{level}}\} \cap \bigcap_{\sigma=0}^{\tau_{\max}} H_{0,\sigma}\right)}{P\left(\bigcap_{\sigma=0}^{\tau_{\max}} H_{0,\sigma}\right)} \leq \sum_{\tau=0}^{\tau_{\max}} \frac{P\left(\{p_\tau < \alpha_{\text{level}}\} \cap H_{0,\tau}\right)}{P\left(\bigcap_{\sigma=0}^{\tau_{\max}} H_{0,\sigma}\right)}$$

$$= \sum_{\tau=0}^{\tau_{\max}} \frac{P(H_{0,\tau})}{P\left(\bigcap_{\sigma=0}^{\tau_{\max}} H_{0,\sigma}\right)} \frac{P\left(\{p_\tau < \alpha_{\text{level}}\} \cap H_{0,\tau}\right)}{P(H_{0,\tau})} = \alpha_{\text{level}} \sum_{\tau=0}^{\tau_{\max}} \frac{P(H_{0,\tau})}{P\left(\bigcap_{\sigma=0}^{\tau_{\max}} H_{0,\sigma}\right)}$$

(3)

705 The last expression has no dependence on the data and depends only on the prior distribution
706 we consider on graphs. Assuming a uniform prior, $P(H_{0,\tau}) = 1/2$. Further,

$$P\left(\bigcap_{\sigma=0}^{\tau_{\max}} H_{0,\sigma}\right) = P(H_{0,0}|H_{0,1} \cap \dots \cap H_{0,\tau_{\max}}) \cdots P(H_{0,\tau_{\max}-1}|H_{0,\tau_{\max}}) \cdot P(H_{0,\tau_{\max}})$$

707 We assume a prior where knowledge of the lack of an edge from X_i to X_j at one lag either
708 increases the probability of lack of an edge between them at other lags or, at least, does not
709 decrease it (independence across lags). Then,

$$P\left(\bigcap_{\sigma=0}^{\tau_{\max}} H_{0,\sigma}\right) \geq P(H_{0,0}) \cdots P(H_{0,\tau_{\max}-1}) \cdot P(H_{0,\tau_{\max}}) = \left(\frac{1}{2}\right)^{\tau_{\max}+1}$$

710 Putting everything together, we get

$$P(\text{Type I Error}) \leq \alpha_{\text{level}} \frac{(\tau_{\max} + 1)/2}{(1/2)^{\tau_{\max}+1}} = \alpha_{\text{level}} (\tau_{\max} + 1) 2^{\tau_{\max}}.$$

711 Note, for analogy, that the correction factor $(\tau_{\max} + 1) 2^{\tau_{\max}}$ takes place of the factor $(\tau_{\max} + 1)$ in
712 a corresponding Bonferroni correction. To have $P(\text{Type I Error})$ less than a prescribed threshold
713 α , we choose

$$\alpha_{\text{level}} = \frac{\alpha}{(\tau_{\max} + 1) 2^{\tau_{\max}}} \quad (4)$$

714 In our experiments with the HCP data, we have $\tau_{\max} = 3$ and $\alpha = 0.01$, giving a per-lag
715 significance threshold of $\alpha_{\text{level}} \simeq 0.0003$. This is notably smaller than the Alpha Level values that
716 maximized F1 scores in simulated Macaque_Full data (0.1 for adjacency F1 score and 0.01 for F1
717 score), and is due to the conservative nature of this correction for temporal multiple comparisons.

718 Computing Functional Graphs

719 In order to calculate the functional graphs for each subject, we consolidated the data from the four
720 sessions of each subject in the HCP and computed the pairwise correlations among all pairs of
721 parcels. To form a binary functional graph, we placed an edge between any two parcels
722 displaying a statistically significant correlation coefficient ($p < 0.01$, t-test for Pearson correlation
723 coefficient).

724 Hyperparameter Selection

725 All the methods we described in Table 1 have at least one main hyperparameter that significantly
726 affects the end result, particularly in terms of edge density. These include ‘PC Alpha’ and ‘Alpha
727 Level’ for PCMCI, ‘PC Alpha’ for PCMCI⁺, ‘Alpha’ for VARLINGAM, DYNOTEARs, and
728 MVGC, and FASK, and ‘BIC Coefficient’ for DGlearn. These hyperparameters were swept over
729 (simultaneously for PCMCI) using the simulated data and selected such that the F1 score with the
730 ground truth graph is maximized in each case. This process was repeated for all algorithms and
731 all experiments (simple graphs, Macaque_SmallDegree, Macaque_Full). Performance metrics
732 such as Recall, Precision, and F1 scores of each method for a range of their hyperparameters are
733 shown in Supplementary Figures 6-12 for the simulated Simple Network graphs, in
734 Supplementary Figures 20-24 for the simulated Macaque_SmallDegree data, and in
735 Supplementary Figures 31-38 for the simulated Macaque_Full data.

736 Time-series algorithms (PCMCI, PCMCI⁺, VARLiNGAM, DYNOTEARs) also have a
737 hyperparameter controlling the number of lags used for causal discovery. Based on our prior
738 work (Nozari et al., 2023), we set this value to 3 for the HCP data (TR = 0.72s), and confirmed its
739 sufficiency based on the contributions of higher-order lags (Figure 6a). For the simulated data,
740 (TR = 1.2s), we used a maximum lag of 2 to match its slower sampling.

741 *Computing F1 Scores, Degrees, and Causal Flows*

742 In our experiments using simulated fMRI data, access to ground truth graphs allows for
743 evaluating the performance of causal discovery methods. In this work, we evaluate causal
744 discovery algorithms as binary classifiers deciding the presence or lack of n^2 edges among n
745 nodes. This allows us to evaluate algorithms using standard classification metrics such as
746 precision, recall, and F1 score (Davis & Goadrich, 2006; Fawcett, 2006; Powers, 2020; Sokolova
747 & Lapalme, 2009; Tharwat, 2020). Given that the F1 score provides a balanced trade-off between
748 precision and recall, we use it as our measure of accuracy. We define two separate metrics, (full)
749 F1 score and adjacency F1 score. For the former, each of the n^2 edges (including any self-loops
750 due to dampening autocorrelation for each node) in the graph is considered as one test sample for
751 classification. In the latter, the ground-truth and learned graphs are first transformed into an
752 undirected graph, placing an edge between two nodes if a directed edge existed in at least one

753 direction. The resulting $\binom{n}{2}$ possible edges are then treated as test samples for classification and
754 computation of adjacency F1 score.

755 To determine the degree and causal flow of a node i in a *binary* directed graph, its in-degree
756 (number of edges pointing toward node i) and out-degree (number of edges originating from node
757 i) are first computed and normalized by the total number of nodes in the graph. The degree of
758 node i is then computed as the sum of the out-degree and in-degree, while the causal flow is
759 obtained by subtracting the in-degree from the out-degree. The same process is followed for
760 weighted graphs except that the calculation of in-degree and out-degree involves a weighted
761 mean. Mathematically,

$$\text{Causal Flow } (i) = \frac{1}{N} \sum_{j=1}^N G(i, j) - \frac{1}{N} \sum_{j=1}^N G(j, i) , \quad i = 1, 2, \dots, N$$
$$\text{Degree } (i) = \frac{1}{N} \sum_{j=1}^N G(i, j) + \frac{1}{N} \sum_{j=1}^N G(j, i) , \quad i = 1, 2, \dots, N$$

762 where G denotes the graph's (binary or weighted) adjacency matrix.

763 **Computing Subnetwork Graphs from Parcel-Level Graphs**

764 Subnetwork graphs were computed by aggregating parcel-level binary graphs into graphs
765 between 16 subnetworks. These subnetworks consist of the standard 7 resting-state
766 subnetworks (Yeo et al., 2011) plus one subcortical subnetwork, separately for the left and right
767 hemispheres. A subnetwork-level graph is then computed for each subject, whereby the weight of
768 an edge from subnetwork i to j is the number of nodes in subnetwork i that connect to nodes in
769 subnetwork j , normalized by the number of all possible edges between these subnetworks. The
770 results are then averaged over the subject, as depicted in Supplementary Figure 51.

771 **Computing**

772 All the computations reported in this study were performed on a Lenovo P620 workstation with
773 AMD 3970X 32-Core processor, Nvidia GeForce RTX 2080 GPU, and 512GB of RAM.

SUPPORTING INFORMATION

774 All the fMRI data used in this work is publicly available. The simulated fMRI benchmarks can be
775 downloaded from <https://github.com/cabal-cmu/Feedback-Discovery> and the
776 human fMRI data can be accessed via the HCP S1200 Release at
777 <https://www.humanconnectome.org/study/hcp-young-adult/document/1200-subjects-data-release>. The Python code for this study is publicly available at
779 https://github.com/nozarilab/2023Arab_CaLLTiF.

AUTHOR CONTRIBUTIONS

780 EN and AG designed and supervised the study; FA performed the research; HJ and MAKP
781 assisted in the analyses of human fMRI data; FA and EN drafted and all authors edited the
782 manuscript.

ACKNOWLEDGMENTS

783 The research conducted in this study was partially supported by NSF award no. 2239654 (to EN),
784 the Canadian Institute for Advanced Research (fellowship awarded to MAKP), and the Air Force
785 Office of Scientific Research under award no. FA9550-20-1-0106 (to MAKP).

COMPETING INTERESTS

786 The authors declare no competing financial interests.

TECHNICAL TERMS

787 **Causal Discovery:** The process of identifying causal relationships between variables from
788 observational data, namely, determining how changing the value of each variable causally
789 influences others.

790 **Contemporaneous Causal Effect:** A causal relationship occurring within the same
791 observation interval when the underlying causal processes are faster than the rate of sampling.

792 **Causal Flow:** The net difference between outgoing and incoming edges in a causal graph,
793 indicating whether a node is a source (positive causal flow) or a sink (negative causal flow).

794 **Partial Correlation:** A statistical measure assessing the direct linear association between two
795 variables after regressing out the effects of a set of conditioned variables.

796 **Conditional Independence Test:** A statistical test that evaluates whether two variables are
797 independent after conditioning on the influence of one or more other variables. Partial correlation
798 is a common method for testing conditional independence in linear models.

799 **Type I Error:** In the context of binary decisions, the incorrect rejection of a true null
800 hypothesis, a.k.a. a false positive.

801 **Correction for Multiple Comparisons:** The process of adjusting statistical tests to reduce the
802 risk of Type I errors when multiple hypotheses are tested simultaneously.

803 **F1 Score:** A balanced performance metric for binary classification, combining precision and
804 recall into a single value that penalizes both false positives and false negatives.

805 **Subnetwork-Level Graph:** A low-dimensional graph in which nodes represent
806 functionally-clustered subnetworks of parcels and edges represent average connectivity between
807 parcels in two subnetworks.

808 **Hemispheric Laterality:** The asymmetry in the distribution of a variable between the two
809 brain hemispheres.

810

811

REFERENCES

812

813 Adams, A., Albin, S., Amunts, K., Asakawa, T., Bernard, A., Bjaalie, J. G., ... others (2020). International brain
814 initiative: an innovative framework for coordinated global brain research efforts. *Neuron*, 105(2), 212–216.

815 Agcaoglu, O., Miller, R., Mayer, A. R., Hugdahl, K., & Calhoun, V. D. (2015). Lateralization of resting state
816 networks and relationship to age and gender. *Neuroimage*, 104, 310–325.

817 Amunts, K., Ebell, C., Muller, J., Telefont, M., Knoll, A., & Lippert, T. (2016). The human brain project: creating
818 a european research infrastructure to decode the human brain. *Neuron*, 92(3), 574–581.

819 Andrews-Hanna, J. R., Smallwood, J., & Spreng, R. N. (2014). The default network and self-generated thought:
820 Component processes, dynamic control, and clinical relevance. *Annals of the New York Academy of Sciences*,
821 1316(1), 29–52.

822 Assaad, C. K., Devijver, E., & Gaussier, E. (2022). Survey and evaluation of causal discovery methods for time
823 series. *Journal of Artificial Intelligence Research*, 73, 767–819.

824 Banks, S. J., Zhuang, X., Bayram, E., Bird, C., Cordes, D., Caldwell, J. Z., ... others (2018). Default mode
825 network lateralization and memory in healthy aging and alzheimer's disease. *Journal of Alzheimer's Disease*,
826 66(3), 1223–1234.

827 Barch, D. M. (2017). Resting-state functional connectivity in the human connectome project: Current status and
828 relevance to understanding psychopathology. *Harvard Review of Psychiatry*, 25(5), 209–217. doi:
829 10.1097/HRP.0000000000000166

830 Barnett, L., & Seth, A. K. (2014). The mvgc multivariate granger causality toolbox: A new approach to
831 granger-causal inference. *Journal of Neuroscience Methods*, 223, 50–68. doi:
832 <https://doi.org/10.1016/j.jneumeth.2013.10.018>

833 Bressler, S. L., & Menon, V. (2010). Large-scale brain networks in cognition: emerging methods and principles.
834 *Trends in Cognitive Sciences*, 14(6), 277–290. Retrieved from
835 <https://www.sciencedirect.com/science/article/pii/S1364661310000896> doi:

836 <https://doi.org/10.1016/j.tics.2010.04.004>

837 Bullmore, E., & Sporns, O. (2012). The economy of brain network organization. *Nature Reviews Neuroscience*,
838 13(5), 336–349.

839 Burgess, G., Kandala, S., Nolan, D., Laumann, T., Power, J., Adeyemo, B., ... Barch, D. (2016). Evaluation of
840 denoising strategies to address motion-correlated artifacts in resting-state functional magnetic resonance imaging
841 data from the human connectome project. *Brain Connectivity*, 6(9), 669–680. doi: 10.1089/brain.2016.0435

842 Buxton, R. B., Wong, E. C., & Frank, L. R. (1998). Dynamics of blood flow and oxygenation changes during brain
843 activation: the balloon model. *Magnetic resonance in medicine*, 39(6), 855–864.

844 Chickering, D. M. (2002a). Learning equivalence classes of bayesian-network structures. *The Journal of Machine
845 Learning Research*, 2, 445–498.

846 Chickering, D. M. (2002b). Optimal structure identification with greedy search. *Journal of machine learning
847 research*, 3(Nov), 507–554.

848 Corbetta, M., Patel, G., & Shulman, G. L. (2008). The reorienting system of the human brain: from environment to
849 theory of mind. *Neuron*, 58(3), 306–324.

850 Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature
851 reviews neuroscience*, 3(3), 201–215.

852 Davis, J., & Goadrich, M. (2006). The relationship between precision-recall and roc curves. In *Proceedings of the
853 23rd international conference on machine learning* (pp. 233–240).

854 Dawid, A. P. (1979). Conditional independence in statistical theory. *Journal of the Royal Statistical Society Series
855 B: Statistical Methodology*, 41(1), 1–15.

856 Ding, M., Chen, Y., & Bressler, S. L. (2006). Granger causality: basic theory and application to neuroscience.
857 *Handbook of time series analysis: recent theoretical developments and applications*, 437–460.

858 Elam, J. (2020). *Hcp data release updates: Known issues and planned fixes*. Retrieved from
859 <https://wiki.humanconnectome.org/docs/HCP%20Data%20Release%20Updates%20Known%20Issues%20and%20Planned%20fixes.html> (Last accessed May 13, 2020)

861 Entner, D., & Hoyer, P. O. (2010). On causal discovery from time series data using fci. *Probabilistic graphical*
862 *models*, 121–128.

863 Essen, D. C. V., Smith, S. M., Barch, D. M., Behrens, T. E., Yacoub, E., & Ugurbil, K. (2013). The wu-minn
864 human connectome project: an overview. *NeuroImage*, 80, 62-79. doi: doi:10.1016/j.neuroimage.2013.05.041

865 Fawcett, T. (2006). An introduction to roc analysis. *Pattern recognition letters*, 27(8), 861–874.

866 Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex.
867 *Cerebral cortex (New York, NY: 1991)*, 1(1), 1–47.

868 Fornito, A., & Bullmore, E. T. (2015). Connectomics: A new paradigm for understanding brain disease. *European*
869 *Neuropsychopharmacology*, 25(5), 733-748. (Advances in neuroscience for improved mental health: an ECNP
870 analysis of progress and prospects) doi: <https://doi.org/10.1016/j.euroneuro.2014.02.011>

871 Friston, K., Harrison, L., & Penny, W. (2003). Dynamic causal modelling. *NeuroImage*, 19(4), 1273-1302. doi:
872 [https://doi.org/10.1016/S1053-8119\(03\)00202-7](https://doi.org/10.1016/S1053-8119(03)00202-7)

873 Friston, K., Moran, R., & Seth, A. K. (2013). Analysing connectivity with granger causality and dynamic causal
874 modelling. *Current opinion in neurobiology*, 23(2), 172–178.

875 Friston, K., Preller, K. H., Mathys, C., Cagnan, H., Heinzle, J., Razi, A., & Zeidman, P. (2019). Dynamic causal
876 modelling revisited. *NeuroImage*, 199, 730-744. doi: <https://doi.org/10.1016/j.neuroimage.2017.02.045>

877 Friston, K. J. (2011). Functional and effective connectivity: A review. *Brain Connectivity*, 1(1), 13–36.

878 Friston, K. J., Kahan, J., Biswal, B., & Razi, A. (2014). A dcm for resting state fmri. *NeuroImage*, 94, 396-407.
879 Retrieved from
880 <https://www.sciencedirect.com/science/article/pii/S1053811913012135> doi:
881 <https://doi.org/10.1016/j.neuroimage.2013.12.009>

882 Gerhardus, A., & Runge, J. (2020). High-recall causal discovery for autocorrelated time series with latent
883 confounders. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, & H. Lin (Eds.), *Advances in neural*
884 *information processing systems* (Vol. 33, pp. 12615–12625). Curran Associates, Inc.

885 Ghassami, A., Yang, A., Kiyavash, N., & Zhang, K. (2020). Characterizing distribution equivalence and structure

886 learning for cyclic and acyclic directed graphs. In *International conference on machine learning* (pp.
887 3494–3504).

888 Glymour, C., Zhang, K., & Spirtes, P. (2019). Review of causal discovery methods based on graphical models.
889 *Frontiers in genetics*, 10, 524.

890 Goebel, R., Roebroeck, A., Kim, D.-S., & Formisano, E. (2003). Investigating directed cortical interactions in
891 time-resolved fmri data using vector autoregressive modeling and granger causality mapping. *Magnetic*
892 *resonance imaging*, 21(10), 1251–1261.

893 Goense, J. B., & Logothetis, N. K. (2008). Neurophysiology of the bold fmri signal in awake monkeys. *Current*
894 *Biology*, 18(9), 631-640. doi: <https://doi.org/10.1016/j.cub.2008.03.054>

895 Granger, C. W. J. (1969). Investigatinag causal relations by econometric models and cross-spectral methods.
896 *Econometrica*, 37(3), 424–438.

897 Greicius, M. D., Krasnow, B., Reiss, A. L., & Menon, V. (2003). Functional connectivity in the resting brain: a
898 network analysis of the default mode hypothesis. *Proceedings of the national academy of sciences*, 100(1),
899 253–258.

900 Hasan, U., Hossain, E., & Gani, M. O. (2023). A survey on causal discovery methods for temporal and
901 non-temporal data. *arXiv preprint arXiv:2303.15027*.

902 Heckerman, D., Geiger, D., & Chickering, D. M. (1995). Learning bayesian networks: The combination of
903 knowledge and statistical data. *Machine learning*, 20, 197–243.

904 Henry, T. R., & Gates, K. M. (2017). Causal search procedures for fmri: review and suggestions. *Behaviormetrika*,
905 44, 193-225.

906 Huettel, S. A., Song, A. W., & McCarthy, G. (2009). *Functional magnetic resonance imaging* (2nd ed.). Sinauer
907 Associates.

908 Hyvärinen, A., Zhang, K., Shimizu, S., & Hoyer, P. O. (2010). Estimation of a structural vector autoregression
909 model using non-gaussianity. *Journal of Machine Learning Research*, 11(56), 1709–1731. Retrieved from
910 <http://jmlr.org/papers/v11/hyvainen10a.html>

911 Jorgenson, L. A., Newsome, W. T., Anderson, D. J., Bargmann, C. I., Brown, E. N., Deisseroth, K., ... others
912 (2015). The brain initiative: developing technology to catalyse neuroscience discovery. *Philosophical*
913 *Transactions of the Royal Society B: Biological Sciences*, 370(1668), 20140164.

914 Lacerda, G., Spirtes, P., Ramsey, J., & Hoyer, P. (2008). Discovering cyclic causal models by independent
915 components analysis. In *Proceedings of the 24th conference on uncertainty in artificial intelligence*.

916 Liao, W., Mantini, D., Zhang, Z., Pan, Z., Ding, J., Gong, Q., ... Chen, H. (2010). Evaluating the effective
917 connectivity of resting state networks using conditional granger causality. *Biological cybernetics*, 102, 57–69.

918 Ljung, L. (1999). System identification: theory for the user. *PTR Prentice Hall, Upper Saddle River, NJ*, 1–14.

919 Logothetis, N. K. (2008). What we can do and what we cannot do with fmri. *Nature*, 453(7197), 869–878.

920 Malinsky, D., & Spirtes, P. (2018). Causal structure learning from multivariate time series in settings with
921 unmeasured confounding. In T. D. Le, K. Zhang, E. Kıcıman, A. Hyvärinen, & L. Liu (Eds.), *Proceedings of*
922 *2018 ACM SIGKDD workshop on causal discovery* (Vol. 92, pp. 23–47). PMLR.

923 Markov, N. T., Ercsey-Ravasz, M. M., Ribeiro Gomes, A., Lamy, C., Magrou, L., Vezoli, J., ... others (2014). A
924 weighted and directed interareal connectivity matrix for macaque cerebral cortex. *Cerebral cortex*, 24(1), 17–36.

925 Meek, C. (1995). Causal inference and causal explanation with background knowledge. In *Conference on*
926 *uncertainty in artificial intelligence*.

927 Meek, C. (1997). *Graphical models: selecting causal and statistical models* (Unpublished doctoral dissertation).
928 Carnegie Mellon University.

929 Mengotti, P., Käsbauer, A.-S., Fink, G. R., & Vossel, S. (2020). Lateralization, functional specialization, and
930 dysfunction of attentional networks. *Cortex*, 132, 206–222.

931 Morton, B. E. (2020). Brain executive laterality and hemisity. *Personality Neuroscience*, 3, e10.

932 Nielsen, J. A., Zielinski, B. A., Ferguson, M. A., Lainhart, J. E., & Anderson, J. S. (2013). An evaluation of the
933 left-brain vs. right-brain hypothesis with resting state functional connectivity magnetic resonance imaging. *PloS*
934 *one*, 8(8), e71275.

935 Nozari, E., Bertolero, M. A., Stiso, J., Caciagli, L., Cornblath, E. J., He, X., ... Bassett, D. S. (2023). Macroscopic
936 resting-state brain dynamics are best described by linear models. *Nature Biomedical Engineering*.

937 Nozari, E., Pasqualetti, F., & Cortés, J. (2019). Heterogeneity of central nodes explains the benefits of time-varying
938 control scheduling in complex dynamical networks. *Journal of Complex Networks*, 7(5), 659–701.

939 Nunez, P. L., & Srinivasan, R. (2006). *Electric fields of the brain: The neurophysics of eeg*. Oxford University
940 Press.

941 Okano, H., Miyawaki, A., & Kasai, K. (2015). Brain/minds: brain-mapping project in japan. *Philosophical
942 Transactions of the Royal Society B: Biological Sciences*, 370(1668), 20140310.

943 Pamfil, R., Sriwattanaworachai, N., Desai, S., Pilgerstorfer, P., Beaumont, P., Georgatzis, K., & Aragam, B. (2020).
944 Dynotears: Structure learning from time-series data. *ArXiv, abs/2002.00498*.

945 Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan kaufmann.

946 Pearl, J. (2009a). Causal inference in statistics: An overview.

947 Pearl, J. (2009b). *Causality: Models, reasoning, and inference*. Cambridge University Press.

948 Poo, M.-m., Du, J.-l., Ip, N. Y., Xiong, Z.-Q., Xu, B., & Tan, T. (2016). China brain project: basic neuroscience,
949 brain diseases, and brain-inspired computing. *Neuron*, 92(3), 591–596.

950 Powers, D. M. (2020). Evaluation: from precision, recall and f-measure to roc, informedness, markedness and
951 correlation. *arXiv preprint arXiv:2010.16061*.

952 Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A
953 default mode of brain function. *Proceedings of the national academy of sciences*, 98(2), 676–682.

954 Ramsey, J. D., Hanson, S. J., Hanson, C., Halchenko, Y. O., Poldrack, R. A., & Glymour, C. (2010). Six problems
955 for causal inference from fmri. *NeuroImage*, 223, 1545–1558. doi:
956 <https://doi.org/10.1016/j.neuroimage.2009.08.065>

957 Ramsey, J. D., Zhang, K., Glymour, M., Romero, R. S., Huang, B., Ebert-Uphoff, I., ... Glymour, C. (2018).
958 Tetrad—a toolbox for causal discovery. In *8th international workshop on climate informatics* (pp. 1–4).

959 Rawls, E., Kummerfeld, E., Mueller, B. A., Ma, S., & Zilverstand, A. (2022). The resting-state causal human
960 connectome is characterized by hub connectivity of executive and attentional networks. *Neuroimage*, 255,
961 119211.

962 Roebroeck, A., Formisano, E., & Goebel, R. (2005). Mapping directed influence over the brain using granger
963 causality and fmri. *Neuroimage*, 25(1), 230–242.

964 Runge, J. (2020). Discovering contemporaneous and lagged causal relations in autocorrelated nonlinear time series
965 datasets. In *Conference on uncertainty in artificial intelligence* (pp. 1388–1397).

966 Runge, J., Nowack, P., Kretschmer, M., Flaxman, S., & Sejdinovic, D. (2019). Detecting and quantifying causal
967 associations in large nonlinear time series datasets. *Science Advances*, 5(11), eaau4996. doi:
968 10.1126/sciadv.aau4996

969 Sanchez-Romero, R., Ramsey, J., Zhang, K., Glymour, M., Huang, B., & Glymour, C. (2019). Estimating
970 feedforward and feedback effective connections from fmri time series: Assessments of statistical methods.
971 *Network Neuroscience*, 86, 274–306.

972 Schaefer, A., Kong, R., Gordon, E. M., Laumann, T. O., Zuo, X.-N., Holmes, A. J., ... Yeo, B. T. T. (2018).
973 Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri. *Cerebral
974 cortex*, 28(9), 3095–3114. doi: 10.1093/cercor/bhx179

975 Seth, A. K., Barrett, A. B., & Barnett, L. (2015). Granger causality analysis in neuroscience and neuroimaging.
976 *Journal of Neuroscience*, 35(8), 3293–3297.

977 Shimizu, S. (2014). Lingam: Non-gaussian methods for estimating causal structures. *Behaviormetrika*, 41, 65–98.

978 Shimizu, S., Hoyer, P. O., Hyvärinen, A., & Kerminen, A. (2006). A linear non-gaussian acyclic model for causal
979 discovery. *Journal of Machine Learning Research*, 7, 2003–2030.

980 Smith, S. M., Miller, K. L., Salimi-Khorshidi, G., Webster, M., Beckmann, C. F., Nichols, T. E., ... Woolrich,
981 M. W. (2011). Network modelling methods for fmri. *NeuroImage*, 875–891. doi:
982 <https://doi.org/10.1016/j.neuroimage.2010.08.063>

983 Sokolova, M., & Lapalme, G. (2009). A systematic analysis of performance measures for classification tasks.
984 *Information processing & management*, 45(4), 427–437.

985 Spirites, P., & Glymour, C. (1991). An algorithm for fast recovery of sparse causal graphs. *Social Science*
986 *Computer Review*, 9(1), 62–72.

987 Spirites, P., Glymour, C. N., & Scheines, R. (2000). *Causation, prediction, and search*. MIT press.

988 Spirites, P., Meek, C., & Richardson, T. (1995). Causal inference in the presence of latent variables and selection
989 bias. In *Proceedings of the eleventh conference on uncertainty in artificial intelligence* (pp. 499–506).

990 Spirites, P., & Zhang, K. (2016). Causal discovery and inference: concepts and recent methodological advances. In
991 *Applied informatics* (Vol. 3, pp. 1–28).

992 Stephan, K. E., & Roebroeck, A. (2012). A short history of causal modeling of fmri data. *Neuroimage*, 62(2),
993 856–863.

994 Sun, X., Liu, G., Poupart, P., & Schulte, O. (2021). Nts-notears: Learning nonparametric temporal dags with
995 time-series data and prior knowledge. *arXiv e-prints*, arXiv–2109.

996 Sutton, S., & Begleiter, H. (1979). Evoked brain potentials and behavior.

997 Sutton, S., Braren, M., Zubin, J., & John, E. (1965). Evoked-potential correlates of stimulus uncertainty. *Science*,
998 150(3700), 1187–1188.

999 Swanson, N., Eichele, T., Pearlson, G., Kiehl, K., Yu, Q., & Calhoun, V. D. (2011). Lateral differences in the
1000 default mode network in healthy controls and patients with schizophrenia. *Human brain mapping*, 32(4),
1001 654–664.

1002 Tharwat, A. (2020). Classification assessment methods. *Applied Computing and Informatics*, 17(1), 168–192.

1003 Tian, Y., Margulies, D. S., Breakspear, M., & Zalesky, A. (2020). Hierarchical organization of the human subcortex
1004 unveiled with functional connectivity gradients. *bioRxiv*. doi: 10.1101/2020.01.22.916357

1005 Tigges, J., Spatz, W., & Tigges, M. (1973). Reciprocal point-to-point connections between parastriate and striate
1006 cortex in the squirrel monkey (saimiri). *Journal of Comparative Neurology*, 148(4), 481–489.

1007 Vossel, S., Geng, J. J., & Fink, G. R. (2014). Dorsal and ventral attention systems: distinct neural circuits but
1008 collaborative roles. *The Neuroscientist*, 20(2), 150–159.

1009 Winder AT, Z. Q., Echagarruga C, & PJ, D. (2017). Weak correlations between hemodynamic signals and ongoing
1010 neural activity during the resting state. *Nature Neuroscience*, 18, 17611769-640. doi:
1011 10.1038/s41593-017-0007-y

1012 Yeo, T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., ... Buckner, R. L. (2011).
1013 The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of*
1014 *neurophysiology*, 106(3), 1125–1165.

1015 Zheng, X., Aragam, B., Ravikumar, P. K., & Xing, E. P. (2018). Dags with no tears: Continuous optimization for
1016 structure learning. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, & R. Garnett (Eds.),
1017 *Advances in neural information processing systems* (Vol. 31). Curran Associates, Inc.