
Causal Discovery for fMRI data: Challenges, Solutions, and a Case Study

Eric Rawls¹

Bryan Andrews¹

Kelvin Lim^{1,2}

Erich Kummerfeld³

¹Psychiatry Behavioral Sciences, University of Minnesota, Minneapolis, Minnesota, USA

²Minneapolis VA Health Care System, Geriatrics Research Education and Clinical Center (GRECC), Minneapolis, Minnesota, USA

³Institute for Health Informatics, University of Minnesota, Minneapolis, Minnesota, USA

Abstract

Designing studies that apply causal discovery requires navigating many researcher degrees of freedom. This complexity is exacerbated when the study involves fMRI data. In this paper we (i) describe nine challenges that occur when applying causal discovery to fMRI data, (ii) discuss the space of decisions that need to be made, (iii) review how a recent case study made those decisions, (iv) and identify existing gaps that could potentially be solved by the development of new methods. Overall, causal discovery is a promising approach for analyzing fMRI data, and multiple successful applications have indicated that it is superior to traditional fMRI functional connectivity methods, but current causal discovery methods for fMRI leave room for improvement.

1 INTRODUCTION: FMRI BRAIN DATA AND EFFECTIVE CONNECTIVITY

Functional Magnetic Resonance Imaging (fMRI) offers the highest-currently-available spatial resolution for three-dimensional images of real-time functional activity across the entire human brain [Glasser et al., 2016]. For this reason, enormous resources have been spent to collect fMRI data from hundreds of thousands of individuals for research purposes [Volkow et al., 2018, Elam et al., 2021, Alfaro-Almagro et al., 2018]. This data is collected and analyzed with the purpose of answering a large variety of scientific questions, such as understanding drivers of adolescent substance use initiation [Volkow et al., 2018].

In this paper we focus on questions related to how activity in different areas of the brain may causally influence activity in other areas of the brain [Friston, 2009]. This can serve a variety of purposes, such as guiding medical inter-

ventions like non-invasive brain stimulation (NIBS) [Horn and Fox, 2020]. For example, deep-brain stimulation targets the subthalamic nucleus (STN), however current NIBS technologies cannot directly manipulate activity in STN [Horn et al., 2017]. The area could potentially be indirectly manipulated through other brain areas, however this requires learning, either for the population or for each individual, which other brain areas have the greatest causal influence on STN. Standard fMRI analysis unfortunately eschews estimating causal effects.

Current standard practice in fMRI analysis is to describe the connections between brain areas with empirical correlations [Biswal et al., 1995]. Partial correlation methods such as *lasso* [Friedman et al., 2008] are also used but are less common [Marrelec et al., 2006]. The connections learned by such methods are called “functional connections”. This term is used to distinguish them from “effective connections” where one area is described as causally influencing another [Friston, 2009, Reid et al., 2019, Pearl, 2000]. Despite the overtly non-causal nature of functional connectivity, it is still used by the larger fMRI research community to identify brain areas that should be targeted with interventions, and more generally to describe how the brain functions. A small but growing community of fMRI researchers are adopting causal discovery analysis (CDA) instead [Spirtes et al., 2000, Camchong et al., 2023, Sanchez-Romero and Cole, 2021, Rawls et al., 2022]. These papers have demonstrated that CDA is capable of recovering information from fMRI data above and beyond what is possible with traditional fMRI connectivity methods, however the details of their approaches vary substantially.

In any project involving CDA, the researcher faces many choices (degrees of freedom). Using CDA requires navigating a wide range of algorithms, properties, assumptions, and settings. Previous fMRI CDA studies have often differed in their specifics, but share a number of common strategies. We focus this paper’s discussion on the decisions made during one project: the development and application of the Greedy Adjacencies with Non-Gaussian Orientations (GANGO)

method to data from the Human Connectome Project (HCP) [Rawls et al., 2022].

The aim of Rawls et al. [2022] was to describe the structure of effective connectivity in the brain commonly found in healthy individuals while at rest, a.k.a. the resting-state causal connectome. This model could then be used to identify causal connectome alterations that may be responsible for psychopathology in people suffering from mental health disorders, as well as to predict the clinical severity of lesions in different brain areas in terms of its impact on the larger causal connectome. To construct this model, however, numerous decisions were made regarding both how to perform the CDA itself and how to clean and process the raw fMRI data. The need for such decisions comes from several challenges that are universal to CDA fMRI studies, which we describe next.

2 PRACTICAL CHALLENGES OF APPLYING CDA TO FMRI DATA

Any research study where CDA is applied to fMRI data will have to confront numerous study design challenges. Here we enumerate 9 challenges that apply to all CDA fMRI studies.

C1: Preprocessing. Raw fMRI data contains numerous artifacts due to a wide variety of physical, biological, and measurement technology factors. For example, there are many artifacts resulting from the fact that the brain is not a rigid, stable object: head motion will obviously impact which locations in space within the scanner correspond to which parts of the brain, but many other less obvious factors such as blinking, swallowing, and changes in blood pressure due to heartbeats, all apply pressure to the brain and causes it to move and change shape. There is a large space of methods for cleaning and preprocessing fMRI data, and these can have a large impact on any fMRI analysis [Parkes et al., 2018, Botvinik-Nezer et al., 2020]. Further, there is an interaction between how the fMRI data is cleaned and preprocessed and what CDA methods are viable. Many popular fMRI preprocessing methods produce Gaussian data, but some popular causal discovery methods require the data to be non-Gaussian [Ramsey et al., 2014]. As such, we can not choose the fMRI preprocessing method and CDA method independently of each other: they must be chosen jointly.

C2: Cycles. Brains are known to contain both positive and negative feedback cycles [Sanchez-Romero et al., 2019, Garrido et al., 2007]. CDA methods that are capable of learning cyclic relationships will thus be preferable, *ceteris paribus*, to CDA methods that can not learn models with cycles. Further, the CDA methods that can accurately learn cycles primarily operate outside the space of Gaussian distributions. However, as already mentioned, many fMRI cleaning meth-




Figure 1: Using CDA on parcellated fMRI data is complicated by the likely presence of both cycles and latent variables. This figure illustrates a toy example of parcels P1, P2, and P3, with unmeasured confounding of P1 and P2 from the peripheral nervous system (PNS). P1 has a cycle of length 1, as activity in P1 directly impacts further activity in P1 (for example, the firing of interneurons in a brain area will suppress the firing of other neurons). There is a 2-cycle between P1 and P2, as each directly causes the other. There is also a 3-cycle between P1, P2, and P3. Many methods capable of learning causal cycles are limited to learning only cycles with length greater than 2, however cycles of all lengths are likely present in brain parcellations.

ods force the data to be Gaussian, thus making it essentially unusable for those methods.

C3: Undersampling. The sampling rate of fMRI imaging is much slower than the rate at which neurons influence each other. That is, typical image acquisitions only sample the brain about every 1-2 seconds [Darányi et al., 2021] and the Blood Oxygen Level Dependent (BOLD) response does not peak for approximately 5-7 seconds following activation of neurons [Buckner, 1998]. Meanwhile, pyramidal neurons can fire up to 10 times per second and interneurons may fire as many as 100 times per second [Csicsvari et al., 1999]. Some recent CDA research has focused on undersampled time series data [Hytinen et al., 2016, 2017, Cook et al., 2017, Solovyeva et al., 2023], however the application of these approaches to parcellated fMRI data remains largely unexplored.

C4: Latents. It is plausible that our measured variables are influenced by some unmeasured common causes, such as haptic or interoceptive feedback from the peripheral nervous system, or even inputs from small brain regions that are not included separately in parcellations but typically lumped together, such as the raphe nucleus (serotonin), the locus coeruleus (norepinephrine), and the ventral tegmental area (dopamine).

C5: Spatial smoothing. fMRI is subject to poorly characterized spatial smoothing resulting from the scanner itself and also from standard preprocessing that typically includes spatial smoothing with a Gaussian kernel [Mikl et al., 2008]. This induces correlations between nearby brain areas. Since these correlations are due only to the measurement and standard preprocessing technologies, they do not reflect causal processes inside the brain. However, current CDA methods will universally attempt to explain these correlations with conjectured causal mechanisms. Performing analysis at the level of parcellations of voxels — biologically interpretable spatially contiguous sets of voxels — is the typical approach that fMRI researchers use to ameliorate this problem.

C6: High Dimensionality. At present, the smallest brain areas that are commonly used for full-brain connectivity analysis are parcellations [Glasser et al., 2016, Schaefer et al., 2018]. These are smaller than most Regions of Interest (ROIs) or other spatially defined “brain networks” (e.g. the “Default Mode Network”), but much larger than voxels. Individual parcels typically comprise 100-200 voxels or more [Glasser et al., 2016]. Voxels are in turn much larger than neurons, containing around 1 million neurons in an 8 mm³ voxel [Ip and Bridge, 2021]. There are multiple different whole-brain parcellations, but they typically include hundreds of parcels. Some examples include the recent multimodal parcellation of the human cortex into 360 parcels by Glasser et al. [2016], and the multiscale parcellation of Schaefer and colleagues that includes between 100 and 1000 parcels [Schaefer et al., 2018]. The analysis algorithm must therefore scale to hundreds or several hundreds of variables, which excludes many CDA methods.

C7: High Density. Brain networks are densely connected. On average, the nodes of parcellated fMRI networks are typically connected to at least 10 other nodes [Rawls et al., 2022]. With some recent exceptions, most CDA methods have reduced performance and greatly increased computational cost on models with such high density [Lam et al., 2022a].

C8: Scale-free structure. Brain networks are scale-free at many different resolutions, including at the resolution of fMRI parcellations [Watts and Strogatz, 1998, Rawls et al., 2022]. This small-world type of connectivity is characterized by having a small number of extremely well-connected nodes, aka hubs. These hubs play critical roles in organizing complex brain functions where multiple regions interact [van den Heuvel and Sporns, 2013, Crossley et al., 2014]. For example, the anterior cingulate cortex (ACC) is densely interconnected with other subcortical and cortical regions, receiving information about emotion and valuation from subcortical brain systems and sending information about the need for control to other brain regions. The extremely high connectivity of these nodes may make them more difficult for some CDA methods to learn, especially as many methods encode sparsity biases that prefer models with more

distributed connectivity (like Erdos-Renyi models) [Erdős et al., 1960, Karoński and Ruciński, 1997].

C9: Limited Samples. fMRI brain data from a single session typically has sample sizes (number of images/frames) that range from several hundred to two thousand [Volkow et al., 2018, Elam et al., 2021, Alfaro-Almagro et al., 2018]. Modern imaging protocols involve imaging at a rate of about 1 capture per second, and the participant is required to stay extremely still. This includes not swallowing and not blinking too much. This is not a comfortable experience, making the total duration of scanning, and thus the total sample size from a single session, necessarily limited. Methods that require more than a few thousand data points are thus not feasible for fMRI, unless they are intended for analyzing multiple sessions or subjects (and thus not focused on our goal of modeling an individual person’s causal connectome at a point in time).

To summarize, the ideal analysis will (1) preprocess the data in a way that cleans as many artifacts as possible while enabling the chosen CDA method, (2) be able to recover causal cycles, (3) be relatively unaffected by or explicitly model temporal undersampling, (4) allow for the possibility of unmeasured confounding, (5) not become biased by spatial smoothing, (6-7) scale to, and retain strong performance on, data with hundreds of densely-connected (average degree 10 or more) variables, (8) retain strong performance for hub nodes in scale free models, (9) achieve all of the above on data with between hundreds and two thousand samples.

3 CDA AND EFFECTIVE CONNECTIVITY METHODS FOR FMRI

When selecting a CDA method for fMRI analysis, the large number of algorithms developed over the last thirty years can be intimidating. However, the majority of CDA methods do not meet the minimum requirements for analyzing parcellated fMRI data. For example, many CDA methods do not scale to problems with hundreds of variables (C6). Moreover, methods that rely on the cross-temporal relationships in the data, such as Granger causality [Granger, 1969, Friston et al., 2014], will not produce meaningful results due to undersampling [Barnett and Seth, 2017] (C3). In practice, researchers usually treat fMRI data as if it were independent and identically distributed while applying CDA methods to avoid this issue. Below, we review popular CDA methods that are appropriate for analyzing parcellated fMRI data. Table 1 compares these methods relative to the 9 challenges.

Greedy Equivalent Search (GES) is a two-phase greedy search algorithm that moves between equivalence classes of Directed Acyclic Graphs (DAGs) by adding or removing conditional independence relations. Fast Greedy Equivalent Search (fGES) is an efficient implementation of GES capa-

	C1	C2	C3	C4	C5	C6	C7	C8	C9
GES	✓	✗	✗	✗	✗	✓	✗	✗	✓
GRaSP	✓	✗	✗	✗	✗	~	✓	✓	✓
Direct LiNGAM	~	✗	✗	✗	✗	✓	✓	✓	✓
GANGO	~	~	✗	✗	✗	✓	~	~	✓
FASk	~	✓	✗	✗	✗	✓	~	~	✓
Two-Step	~	✓	✗	~	✗	✓	~	~	✓

Table 1: A comparison of the discussed CDA methods relative to the 9 challenges where (✓) denotes the challenge is addressed, (~) denotes the challenge is partially addressed, and (✗) denotes the challenge is not addressed. Note that while the same challenge may be addressed by multiple algorithms, it is still possible that some algorithms address have better performance on this challenge than others. For example, GES, GRaSP, Direct LiNGAM, and GANGO appear to have better accuracy on low samples (C9) than FASk and Two-Step.

ble of scaling up to a million variables (C6). It does not have any special preprocessing requirements (C1), and has good model performance on limited samples (C9). However, this algorithm is intended for sparsely connected networks and suffers in terms of performance and scalability on densely connected problems (C7) [Chickering, 2002, Ramsey et al., 2017].

Greedy Relaxations of the Sparsest Permutation (GRaSP) is a hierarchy of greedy search algorithms that move between variable orderings. For brevity, we will use the GRaSP acronym to refer to the most general algorithm in the hierarchy. Starting from a random order, GRaSP repeatedly iterates over pairs of variables that are adjacent in the DAG constructed by applying the Grow Shrink (GS) [Margaritis and Thrun, 1999] in a manner consistent with the order, modifying the order in a way consistent with flipping the corresponding edge in the DAG. GRaSP does not have any special preprocessing requirements (C1), and has good model performance on limited samples (C9). However, unlike fGES, GRaSP retains its good performance on high density models (C7) but only scales to one or two hundred variables (C6) [Lam et al., 2022b].

Direct Linear Non-Gaussian Acyclic Model (LiNGAM) is a greedy algorithm that constructs a variable ordering based on a pairwise orientation criterion [Hyvärinen and Smith, 2013]. Starting from an empty list, the order is constructed by adding one variable at a time, the one that maximizes the pairwise orientation criterion, until a full ordering is constructed. Once the order is constructed, it is projected to a DAG. The pairwise orientation criterion uses measures non-Gaussianity so some preprocessing techniques cannot be used (C1). Moreover, LiNGAM cannot learn cycles (C2). That being said, the method scales fairly well and has not problem learning scale-free structures (C6), (C7), and (C9) [Shimizu et al., 2011].

The next three methods broadly fall into the same class of algorithms. These approaches are two stepped approaches where the first step learns a graph using an existing CDA method and the second step augments and (re)orients that

edges of the graph learned in the first step; this general approach was pioneered by Hoyer et al. [2008]. All three of these methods use non-Gaussianity and thus some preprocessing techniques cannot be used (C1). Moreover, the properties of these methods are impacted their algorithm choice in the first step, for example, if the chosen fails to scale in some aspect, then so will the overall method (C6 - C9).

Greedy Adjacencies with Non-Gaussian Orientations (GANGO) uses fGES [Ramsey et al., 2017] as a first step in order to learn the adjacencies and then uses the RSkew pairwise orientation rule [Hyvärinen and Smith, 2013], also referred to as robust skew, for orientations. This method allows for cycles (with the exception of two-cycles) and scales well [Rawls et al., 2022].

Fast Adjacency Skewness (FASk) uses fast adjacency search (FAS), which is the adjacency phase of the PC algorithm, as a first step in order to learn the adjacencies and then uses a series of tests to add additional adjacencies, orient two-cycles, and orient directed edges. This method allows for cycles and scales well [Sanchez-Romero et al., 2019].

Two-Step uses adaptive lasso or FAS as a first step in order to learn the adjacencies and then uses independent subspace analysis (ISA) or independent component analysis (ICA) if no latent confounders are identified. This method allows for cycles and latent confounding, and scales well [Sanchez-Romero et al., 2019].

4 INTERDEPENDENCIES AMONG CDA AND FMRI PROCESSING METHODS

This section briefly covers some of the more important ways in which choice of CDA method and choice of fMRI processing methods interact, and how these complexities can be successfully navigated.

First, as discussed above, the better CDA methods for learning cycles reliably requires using non-Gaussian statistics,

however many preprocessing methods force the fMRI data to conform to a Gaussian distribution. While the Cyclic Causal Discovery (CCD) algorithm [Richardson, 1996] can recover causal graphs with cycles from Gaussian data, this algorithm performs poorly on finite samples and is rarely used. Methods exploiting non-Gaussian structure in BOLD data achieve higher precision and recall with simulated BOLD data [19]. Fortunately, there are approaches to preprocessing fMRI that do not completely remove the non-Gaussian signal. Those approaches are thus recommended for using CDA on fMRI data.

Preprocessing removes artifacts and recovers physiological brain signals via some combination of temporal filtering, spatial smoothing, independent components analysis (ICA), and confound regression [Glasser et al., 2013]. Some of these steps, particularly temporal filtering, can drastically modify the data distribution. For example, [Ramsey et al., 2014] demonstrated that certain high-pass temporal filters made parcellated fMRI time series more Gaussian. This effect was particularly strong for Butterworth filters, which were applied in [Smith et al., 2011]. As such, it is likely that the results of [Smith et al., 2011] were unrealistically pessimistic with regards to methods that assume non-Gaussianity. This effect holds true for the filter built into the Statistical Parametric Mapping (SPM) software, while being negligible for the filter built into the fMRIB Software Library (FSL) software. Thus, high-pass filtering including the specific software and filter used is a critical point of attention during data preparation for CDA.

For filtered data that maintain non-Gaussianity in the BOLD signal, it's crucial to confirm the data meet distributional assumptions. A recent cortex-wide human causal connectome analysis discovered that minimally preprocessed cortical BOLD signal was non-Gaussian for all subjects [Rawls et al., 2022]. In that same dataset, however, the subcortical parcel time series are not non-Gaussian. This could stem from Gaussian noise corrupting non-Gaussian BOLD activity, especially since subcortical regions typically display low signal-to-noise ratios. There is potential for enhancing BOLD data's compatibility with CDA methods by employing newer techniques that eliminate Gaussian noise, like the NORDIC method [Moeller et al., 2021, Vizioli et al., 2021], which suppresses Gaussian thermal noise from high-resolution scan parameters. However, to date we are unaware of any studies that pursue this combination of preprocessing methods and CDA.

5 ADDITIONAL COMPLICATIONS OF CDA ON FMRI

Regarding the challenges of high dimensionality and limited samples, this area fortunately has numerous causal discovery solutions. All of the methods discussed in Section 3 are capable of scaling to the number of parcels found in the

most widely used parcellations, while maintaining good performance (although both runtime and performance can vary substantially across these methods).

The structural challenges of high-density and scale-free models also have some solutions. In particular, recently developed permutation-based methods such as GRaSP and BOSS both retain their high performance as model density increases. These methods have increased computational cost compared to faster methods like fGES, but can still scale comfortably to hundreds of parcels, even on a personal computer. Most other methods appear to have substantial drops in performance as density increases, so using one of the few methods that tolerates high density models is recommended.

Since fMRI preprocessing may leave non-Gaussian marginal distributions, it's worth considering whether the CDA methods that assume a linear-Gaussian model still perform well. In general, such methods retain their performance for edge adjacencies [Smith et al., 2011], while their performance on edge orientations has mixed results. There are known cases where the orientations become essentially random [Smith et al., 2011], while we have observed other cases where the orientations only exhibit a slight drop in accuracy.

Collectively, the challenges and available tools point towards a particular approach:

1. to ensure sample size is not too small (C9), analyze data from study protocols that allow for adequate scanner time for each individual session;
2. for preprocessing (C1), remove as many fMRI artifacts as possible while retaining as much non-Gaussianity in the marginals of the parcellated time-series data as possible;
3. for undersampling (C3) and spatial smoothing (C5), use a cross-sectional approach to analyze parcellations;
4. due to high dimensionality (C6), high density (C7), and scale-free brain structure (C8), use a scalable high-density-tolerant method to learn the skeleton (adjacencies) of the parcels;
5. in order to learn cycles (C2) use a non-Gaussian orientation method to re-orient edges;
6. In consideration of possible latent confounding (C4), one can either
 - (a) use a scalable CDA method capable of learning both cycles and latent confounding, or
 - (b) focus the primary results on aggregate features of the model, such as connections among multiple parcellations in shared networks, rather than on individual edges.

The next section reviews a project where this general approach was taken.

6 CASE STUDY: APPLICATION OF THIS APPROACH TO THE HUMAN CONNECTOME PROJECT (HCP)

A previous project can serve as an example of the above thought process [Rawls et al., 2022]. In that project, the authors made the following considerations with respect to the 9 challenges.

Challenge 1 (preprocessing): The minimal HCP processing pipeline [Glasser et al., 2013] was used, to conserve as much non-Gaussian signal as possible and enable the use of non-Gaussian CDA methods for learning cycles. Non-Gaussianity of data were statistically verified by simulating surrogate Gaussian data for comparison.

Challenge 2 (cycles): RSkew was used to re-orient edges [Hyvärinen and Smith, 2013] after confirming non-Gaussianity of the preprocessed cortical parcellations, enabling discovery of cycles involving three or more variables.

Challenge 3 (undersampling): The time series element of the data was ignored, and the parcellated time series were instead analyzed as cross-sectional data. While this approach does not make use of the available time-order information, it avoids relying on the heavily undersampled time dimension of the data.

Challenge 4 (latents): No effort was made to directly model or account for latent variables. Findings were reported at an aggregate level rather than individual edges.

Challenge 5 (spatial smoothing): These data were only minimally smoothed (2 mm in surface space) [Glasser et al., 2013], thus excessive smoothing was not introduced in the data. In addition, the parcellated time series was analyzed rather than voxels, which further reduced the impact of smoothing.

Challenge 6 (high dimensionality): A 360-node cortical parcellation was used [Glasser et al., 2016]. FGES was used for the more computationally intensive adjacency search, which is among the most scalable methods for CDA [Ramsey et al., 2017].

Challenge 7 (high density): The study used fGES, which scales well but can have lower performance for extremely dense graphs. Better performance for dense brain graphs could potentially be achieved by applying a high-density-tolerant algorithm such as GRaSP [Lam et al., 2022a].

Challenge 8 (scale free): Rawls et al. [2022] reported the existence of nodes that were more highly connected than expected under chance, which is characteristic of scale-free networks. See Figure 3. However, for especially highly-connected hub regions, some methods such as GRaSP might provide higher precision for assessing scale-free structure in future studies.

Challenge 9 (limited samples): The HCP collected two

runs of fMRI per day of 1200 images per run. The study applied CDA to concatenated standardized time series from these two runs, thus the number of samples was extremely high (2400 total). The CDA methods that were selected are also known to have good finite-sample performance.

Overall: This recent large-scale application of CDA for deriving individualized causal connectomes addressed many of the challenges we identified. However, the challenges of high density and scale-free connectivity could potentially be better addressed by applying newer permutation-based CDA methods [Lam et al., 2022b]. Several challenges, such as limited samples, spatial smoothing, and preprocessing, were partially or entirely solved by the specific data set the method was applied to, and might pose problems in other data sets.

6.1 RESULTS FROM THE HCP CASE STUDY

Here we briefly review the findings from the HCP case study by Rawls et al. [2022]. This study developed the GANGO causal connectivity method, which was applied to $n=442$ resting-state fMRI data sets. The connectomes produced were extremely sparse (2.25% edge density) compared to Pearson correlation connectomes, which are often thresholded to an edge density of 5-50%. Nevertheless, graphs produced by GANGO were fully connected in nearly all cases, which was not the case for standard Pearson correlation graphs.

The GANGO method produced graphs with a scale-free degree distribution. More specifically, the degree distributions were skewed by existence of hub nodes with very high connectivity (some with total degree exceeding 20). See Figure 3. These hub nodes were disproportionately concentrated in brain networks tied to attention and executive control, while Pearson correlations instead emphasized hub connectivity of early sensory regions. Graphs produced by the GANGO method also show small-world connectivity, with global efficiency nearly as high as random graphs but local efficiency much higher than random graphs. Overall, this case study showed that a causal discovery algorithm specifically designed to meet the unique challenges of fMRI data recovers physiologically plausible connectomes with small-world and scale-free connectivity patterns characteristic of biological networks.

7 RESEARCH GAPS AND PROMISING FUTURE DIRECTIONS

We have outlined nine challenges researchers will face when attempting to apply CDA to parcellations of fMRI data, as well as some available CDA technologies and their ability to overcome those challenges. The case study discussed in Section 6 attempted to use a mixture of strategies to over-

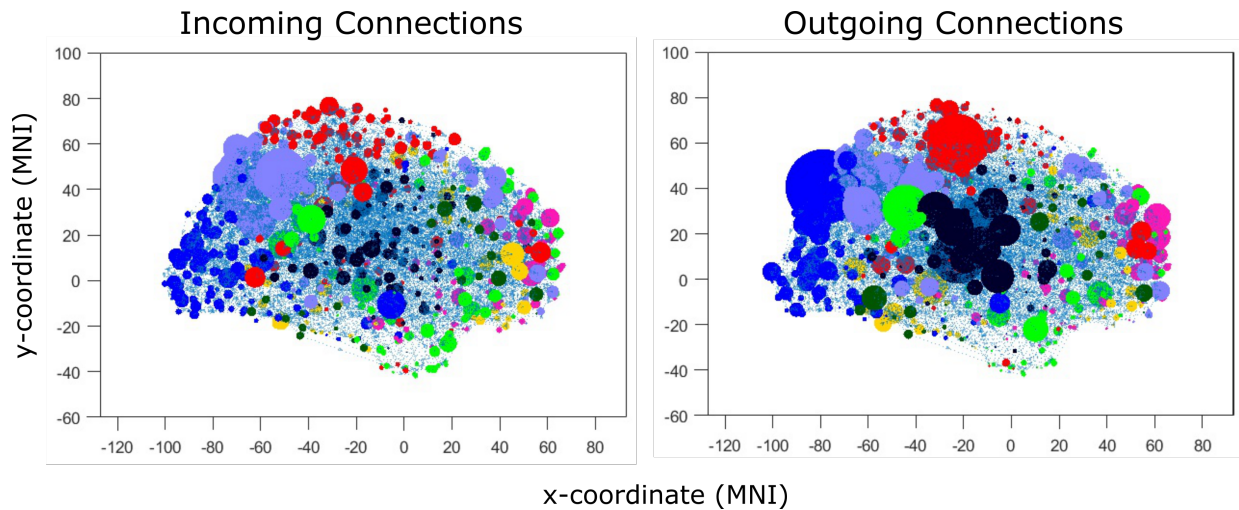


Figure 2: Application of CDA methods to BOLD data allows estimating patterns of directed connectivity, separating incoming and outgoing connections. In these brain images, circles are plotted at the 3D centroids of cortical parcels and the size of the circle indicates the number of connections (the degree) of the parcel. Individual causal connectomes reveal a striking dissociation between brain regions with especially high indegree and outdegree. Different colors refer to different brain functional communities (resting-state networks).

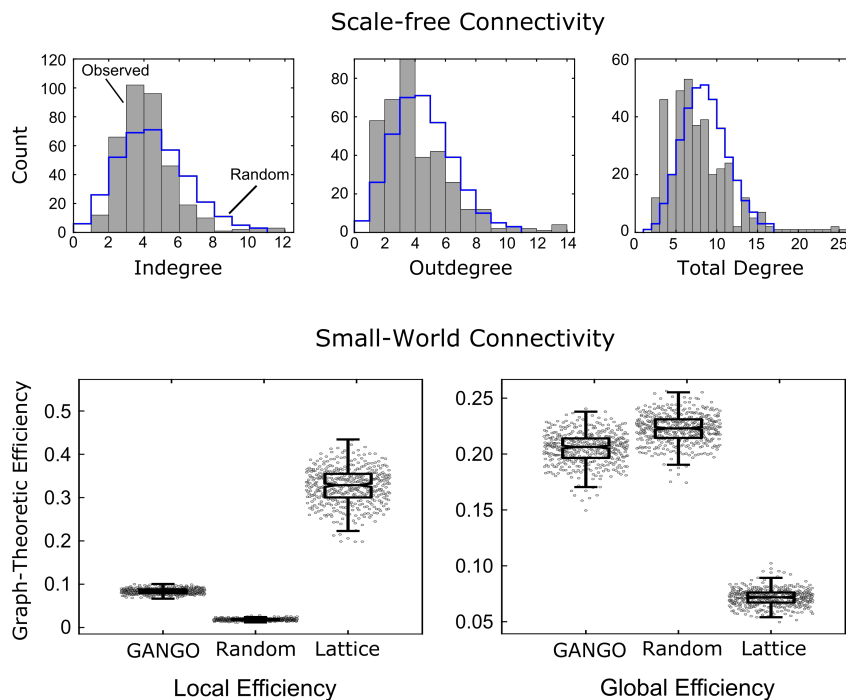


Figure 3: The case study that developed the GANGO framework found that the causal connectomes produced were both scale-free and small-world graphs. The scale-free nature of GANGO connectomes is demonstrated by degree distributions that were more skewed than random graphs. The small-world nature of GANGO connectomes is demonstrated by global efficiency that is nearly equal to random graphs but higher than lattice graphs, combined with local efficiency that is greater than random graphs. Panels illustrating scale-free connectivity are adapted with permission from Rawls et al. [2022].

come those challenges, but many challenges remained only partially addressed, or were even largely ignored. In this section, we review the current research gaps as we perceive them, and point towards future directions to empower future applications of CDA to better elucidate the brain’s causal connectome for both scientific and medical purposes.

Gap 1: CDA methods for high-dimensional, high-density, scale-free models. Previous work [Lam et al., 2022a] has shown that many popular CDA methods unfortunately do not perform well when nodes have larger numbers of connections to other nodes. The only currently published method that appears to overcome this limitation is limited to about 100 variables, which is substantially fewer than most parcellations [Glasser et al., 2016, Schaefer et al., 2018]. We are aware of research on new algorithms and CDA implementation technologies that may overcome this gap, however that work has not yet been published. However, such methods could be used as a replacement for methods like fGES in future fMRI studies.

Gap 2: Reliance on skewed data. While it is the case for some fMRI data that minimal preprocessing is able to retain statistically significant skew, we are also aware of other fMRI data where even after only minimal preprocessing the data are not significantly skewed. One possible future direction would be to incorporate additional information from higher moments such as kurtosis to ensure that the non-Gaussian orientation methods can be used as much as possible. Another approach would be to distill as much non-Gaussian signal as possible using a method like independent components analysis (ICA) [Comon, 1994] to construct a new set of features from the parcellated time series, and then perform CDA on the maximally non-Gaussian components of each parcel instead.

Gap 3: Latent variables. While there exist CDA methods that do not assume causal sufficiency, and thus can tolerate and even identify unmeasured common causes, they generally have significant difficulties with other challenges. For example, the standard methods for handling latent variables, like FCI and GFCA, do not allow for cycles, have limited scalability, and perform poorly for high-density models. The Two-Step algorithm [Sanchez-Romero et al., 2019] can in theory incorporate unmeasured confounding in its models, however we are not aware of any theoretical or practical evaluation of its performance in the presence of unmeasured confounding.

Gap 4: Extension to other brain imaging technologies. Future exploration should expand CDA methodology to neural temporal data beyond BOLD signals (fMRI data). For example, electroencephalography (EEG) provides a dynamic view of brain activation with exceptional temporal resolution. Current EEG causal connectivity analysis techniques, such as Granger causality, are fruitful yet limited. These current methods do not differentiate brain oscillations

from aperiodic activity, which is critical given recent evidence that aperiodic activity sometimes wholly explains group differences in power spectral density [Merkin et al., 2023]. Techniques have recently emerged separating aperiodic and oscillatory contributions to EEG power spectra [Donoghue et al., 2020], even extending to time-frequency domain for time-resolved separation [Wilson et al., 2022]. However, these haven’t been incorporated into neural connectivity analyses. EEG data also present challenges in identifying effective connectivity patterns due to volume conduction – instant, passive electricity conduction through the brain separate from actual neural interactions, resulting in non-independent EEG sensor-level estimates [Nunez and Srinivasan, 2006]. This non-independence hampers CDA, necessitating volume conduction removal from data before connectivity estimation (EEG preprocessing). We suggest these obstacles could be mitigated by first removing volume conduction from EEG data, then separating aperiodic and oscillatory spectral contributions. Applying CDA to isolated EEG oscillatory power estimates could potentially reveal effective connectivity patterns unhindered by aperiodic activity.

8 CONCLUSION

Here we have outlined nine challenges that will be faced by researchers attempting to use CDA for fMRI effective connectivity analysis, and presented a recent case study that attempted to resolve at least some of these challenges. We have also discussed challenges that remain following this case study, such as the continuing search for CDA methods that can discover densely connected graphs and hub nodes with especially high connectivity resulting from scale-free connectivity profiles.

In summary, there are a number of decisions faced by researchers who hope to use CDA for fMRI analysis. By openly discussing these researcher degrees-of-freedom and a recent attempt to resolve these decisions, we hope to foster continued interest in and engagement with the idea that CDA can provide a data-driven method for reconstructing human causal connectomes.

Author Contributions

All authors contributed to the ideas in this paper and contributed to its drafting. All authors have reviewed and edited the paper’s content for correctness.

Acknowledgements

ER was supported by the National Institutes of Health’s National Center for Advancing Translational Sciences, grants TL1R002493 and UL1TR002494. BA was supported by the

Comorbidity: Substance Use Disorders and Other Psychiatric Conditions Training Program T32DA037183. EK was supported by grants P50 MH119569 and UL1TR002494.

References

- Fidel Alfaro-Almagro, Mark Jenkinson, Neal K Bangerter, Jesper L R Andersson, Ludovica Griffanti, Gwenaëlle Douaud, Stamatios N Sotiropoulos, Saad Jbabdi, Moises Hernandez-Fernandez, Emmanuel Vallee, Diego Vidaurre, Matthew Webster, Paul McCarthy, Christopher Rorden, Alessandro Daducci, Daniel C Alexander, Hui Zhang, Iulius Dragonu, Paul M Matthews, Karla L Miller, and Stephen M Smith. Image processing and quality control for the first 10,000 brain imaging datasets from UK biobank. *Neuroimage*, 166:400–424, February 2018.
- Lionel Barnett and Anil K Seth. Detectability of granger causality for subsampled continuous-time neurophysiological processes. *Journal of neuroscience methods*, 275:93–121, 2017.
- B Biswal, F Z Yetkin, V M Haughton, and J S Hyde. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn. Reson. Med.*, 34(4):537–541, October 1995.
- Rotem Botvinik-Nezer, Felix Holzmeister, Colin F Camerer, Anna Dreber, Juergen Huber, Magnus Johannesson, Michael Kirchler, Roni Iwanir, Jeanette A Mumford, R Alison Adcock, Paolo Avesani, Blazej M Baczowski, Aahana Bajracharya, Leah Bakst, Sheryl Ball, Marco Barilari, Nadège Bault, Derek Beaton, Julia Beitner, Roland G Benoit, Ruud M W J Berkers, Jamil P Bhanji, Bharat B Biswal, Sebastian Bobadilla-Suarez, Tiago Bortolini, Katherine L Bottenhorn, Alexander Bowring, Senne Braem, Hayley R Brooks, Emily G Brudner, Cristian B Calderon, Julia A Camilleri, Jaime J Castellon, Luca Cecchetti, Edna C Cieslik, Zachary J Cole, Olivier Collignon, Robert W Cox, William A Cunningham, Stefan Czoschke, Kamalaker Dadi, Charles P Davis, Alberto De Luca, Mauricio R Delgado, Lysia Demetriou, Jeffrey B Dennison, Xin Di, Erin W Dickie, Ekaterina Dobryakova, Claire L Donnat, Juergen Dukart, Niall W Duncan, Joke Durnez, Amr Eed, Simon B Eickhoff, Andrew Erhart, Laura Fontanesi, G Matthew Fricke, Shiguang Fu, Adriana Galván, Remi Gau, Sarah Genon, Tristan Glatard, Enrico Glerean, Jelle J Goeman, Sergej A E Golowin, Carlos González-García, Krzysztof J Gorgolewski, Cheryl L Grady, Mikella A Green, João F Guassi Moreira, Olivia Guest, Shabnam Hakimi, J Paul Hamilton, Roeland Hancock, Giacomo Handjaras, Bronson B Harry, Colin Hawco, Peer Herholz, Gabrielle Herman, Stephan Heunis, Felix Hoffstaedter, Jeremy Hogeveen, Susan Holmes, Chuan-Peng Hu, Scott A Huettel, Matthew E Hughes, Vittorio Iacovella, Alexandru D Iordan, Peder M Isager, Ayse I Isik, Andrew Jahn, Matthew R Johnson, Tom Johnstone, Michael J E Joseph, Anthony C Juliano, Joseph W Kable, Michalis Kassinopoulos, Cemal Koba, Xiang-Zhen Kong, Timothy R Kosciak, Nuri Erkut Kucukboyaci, Brice A Kuhl, Sebastian Kupek, Angela R Laird, Claus Lamm, Robert Langner, Nina Lauharatanahirun, Hongmi Lee, Sangil Lee, Alexander Leemans, Andrea Leo, Elise Lesage, Flora Li, Monica Y C Li, Phui Cheng Lim, Evan N Lintz, Schuyler W Liphardt, Annabel B Losecaat Vermeer, Bradley C Love, Michael L Mack, Norberto Malpica, Theo Marins, Camille Maumet, Kelsey McDonald, Joseph T McGuire, Helena Melero, Adriana S Méndez Leal, Benjamin Meyer, Kristin N Meyer, Glad Mihai, Georgios D Mitsis, Jorge Moll, Dylan M Nielson, Gustav Nilsson, Michael P Notter, Emanuele Olivetti, Adrian I Onicas, Paolo Papale, Kaustubh R Patil, Jonathan E Peelle, Alexandre Pérez, Doris Pischetta, Jean-Baptiste Poline, Yanina Prystauka, Shruti Ray, Patricia A Reuter-Lorenz, Richard C Reynolds, Emiliano Ricciardi, Jenny R Rieck, Anais M Rodriguez-Thompson, Anthony Romy, Taylor Salo, Gregory R Samanez-Larkin, Emilio Sanz-Morales, Margaret L Schlichting, Douglas H Schultz, Qiang Shen, Margaret A Sheridan, Jennifer A Silvers, Kenny Skagerlund, Alec Smith, David V Smith, Peter Sokol-Hessner, Simon R Steinkamp, Sarah M Tashjian, Bertrand Thirion, John N Thorp, Gustav Tinghög, Loreen Tisdall, Steven H Tompson, Claudio Toro-Serey, Juan Jesus Torre Tresols, Leonardo Tozzi, Vuong Truong, Luca Turella, Anna E van 't Veer, Tom Verguts, Jean M Vettel, Sagana Vijayarajah, Khoi Vo, Matthew B Wall, Wouter D Weeda, Susanne Weis, David J White, David Wisniewski, Alba Xifra-Porxas, Emily A Yearling, Sangsuk Yoon, Rui Yuan, Kenneth S L Yuen, Lei Zhang, Xu Zhang, Joshua E Zosky, Thomas E Nichols, Russell A Poldrack, and Tom Schonberg. Variability in the analysis of a single neuroimaging dataset by many teams. *Nature*, 582(7810):84–88, June 2020.
- R L Buckner. Event-related fMRI and the hemodynamic response. *Hum. Brain Mapp.*, 6(5-6):373–377, 1998.
- Jazmin Camchong, Donovan Roediger, Mark Fiecas, Casey S Gilmore, Matt Kushner, Erich Kummerfeld, Bryon A Mueller, and Kelvin O Lim. Frontal tDCS reduces alcohol relapse rates by increasing connections from left dorsolateral prefrontal cortex to addiction networks. *Brain Stimul.*, June 2023.
- David Maxwell Chickering. Optimal structure identification with greedy search. *Journal of machine learning research*, 3(Nov):507–554, 2002.
- Pierre Comon. Independent component analysis, a new concept? *Signal processing*, 36(3):287–314, 1994.
- John W Cook, David Danks, and Sergey M Plis. Learning dynamic structure from undersampled data. In *Proceedings of the UAI Causality Workshop*, 2017.

- Nicolas A Crossley, Andrea Mechelli, Jessica Scott, Francesco Carletti, Peter T Fox, Philip McGuire, and Edward T Bullmore. The hubs of the human connectome are generally implicated in the anatomy of brain disorders. *Brain*, 137(Pt 8):2382–2395, August 2014.
- J Csicsvari, H Hirase, A Czurkó, A Mamiya, and G Buzsáki. Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving rat. *J. Neurosci.*, 19(1): 274–287, January 1999.
- Virág Darányi, Petra Hermann, István Homolya, Zoltán Vidnyánszky, and Zoltan Nagy. An empirical investigation of the benefit of increasing the temporal resolution of task-evoked fMRI data with multi-band imaging. *MAGMA*, 34(5):667–676, October 2021.
- Thomas Donoghue, Matar Haller, Erik J Peterson, Paroma Varma, Priyadarshini Sebastian, Richard Gao, Torben Noto, Antonio H Lara, Joni D Wallis, Robert T Knight, et al. Parameterizing neural power spectra into periodic and aperiodic components. *Nature neuroscience*, 23(12): 1655–1665, 2020.
- Jennifer Stine Elam, Matthew F Glasser, Michael P Harms, Stamatios N Sotiropoulos, Jesper L R Andersson, Gregory C Burgess, Sandra W Curtiss, Robert Oostenveld, Linda J Larson-Prior, Jan-Mathijs Schoffelen, Michael R Hodge, Eileen A Cler, Daniel M Marcus, Deanna M Barch, Essa Yacoub, Stephen M Smith, Kamil Ugurbil, and David C Van Essen. The human connectome project: A retrospective. *Neuroimage*, 244:118543, December 2021.
- Paul Erdős, Alfréd Rényi, and Others. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci.*, 5(1): 17–60, 1960.
- Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, July 2008.
- Karl Friston. Causal modelling and brain connectivity in functional magnetic resonance imaging. *PLoS Biol.*, 7(2): e33, February 2009.
- Karl J Friston, André M Bastos, Ashwini Oswal, Bernadette van Wijk, Craig Richter, and Vladimir Litvak. Granger causality revisited. *Neuroimage*, 101:796–808, 2014.
- Marta I Garrido, James M Kilner, Stefan J Kiebel, and Karl J Friston. Evoked brain responses are generated by feedback loops. *Proc. Natl. Acad. Sci. U. S. A.*, 104(52): 20961–20966, December 2007.
- Matthew F Glasser, Stamatios N Sotiropoulos, J Anthony Wilson, Timothy S Coalson, Bruce Fischl, Jesper L Andersson, Junqian Xu, Saad Jbabdi, Matthew Webster, Jonathan R Polimeni, David C Van Essen, Mark Jenkinson, and WU-Minn HCP Consortium. The minimal pre-processing pipelines for the human connectome project. *Neuroimage*, 80:105–124, October 2013.
- Matthew F Glasser, Timothy S Coalson, Emma C Robinson, Carl D Hacker, John Harwell, Essa Yacoub, Kamil Ugurbil, Jesper Andersson, Christian F Beckmann, Mark Jenkinson, Stephen M Smith, and David C Van Essen. A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615):171–178, August 2016.
- Clive WJ Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society*, pages 424–438, 1969.
- Andreas Horn and Michael D Fox. Opportunities of connectomic neuromodulation. *Neuroimage*, 221:117180, November 2020.
- Andreas Horn, Martin Reich, Johannes Vorwerk, Ningfei Li, Gregor Wenzel, Qianqian Fang, Tanja Schmitz-Hübsch, Robert Nickl, Andreas Kupsch, Jens Volkmann, Andrea A Kühn, and Michael D Fox. Connectivity predicts deep brain stimulation outcome in parkinson disease. *Ann. Neurol.*, 82(1):67–78, July 2017.
- Patrik O Hoyer, Aapo Hyvärinen, Richard Scheines, Peter Spirtes, Joseph Ramsey, Gustavo Lacerda, and Shohei Shimizu. Causal discovery of linear acyclic models with arbitrary distributions. In *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence*, pages 282–289, 2008.
- Antti Hyttinen, Sergey Plis, Matti Järvisalo, Frederick Eberhardt, and David Danks. Causal discovery from subsampled time series data by constraint optimization. In Alessandro Antonucci, Giorgio Corani, and Campos, editors, *Proceedings of the Eighth International Conference on Probabilistic Graphical Models*, volume 52 of *Proceedings of Machine Learning Research*, pages 216–227, Lugano, Switzerland, 2016. PMLR.
- Antti Hyttinen, Sergey Plis, Matti Järvisalo, Frederick Eberhardt, and David Danks. A constraint optimization approach to causal discovery from subsampled time series data. *Int. J. Approx. Reason.*, 90:208–225, November 2017.
- Aapo Hyvärinen and Stephen M Smith. Pairwise likelihood ratios for estimation of Non-Gaussian structural equation models. *J. Mach. Learn. Res.*, 14(Jan):111–152, January 2013.
- Betina Ip and Holly Bridge. Mapping the visual world to the human brain. *Elife*, 10, December 2021.
- Micha Karoński and Andrzej Ruciński. The origins of the theory of random graphs. In Ronald L Graham and

- Jaroslav Nešetřil, editors, *The Mathematics of Paul Erdős I*, pages 311–336. Springer Berlin Heidelberg, Berlin, Heidelberg, 1997.
- Wai-Yin Lam, Bryan Andrews, and Joseph Ramsey. Greedy relaxations of the sparsest permutation algorithm. In James Cussens and Kun Zhang, editors, *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence*, volume 180 of *Proceedings of Machine Learning Research*, pages 1052–1062. PMLR, 2022a.
- Wai-Yin Lam, Bryan Andrews, and Joseph Ramsey. Greedy relaxations of the sparsest permutation algorithm. In *Uncertainty in Artificial Intelligence*, pages 1052–1062. PMLR, 2022b.
- Dimitris Margaritis and Sebastian Thrun. Bayesian network induction via local neighborhoods. In S. Solla, T. Leen, and K. Müller, editors, *Advances in Neural Information Processing Systems*, volume 12. MIT Press, 1999.
- Guillaume Marrelec, Alexandre Krainik, Hugues Dufau, Mélanie Péligrini-Issac, Stéphane Lehericy, Julien Doyon, and Habib Benali. Partial correlation for functional brain interactivity investigation in functional MRI. *Neuroimage*, 32(1):228–237, August 2006.
- Ashley Merkin, Sabrina Sghirripa, Lynton Graetz, Ashleigh E Smith, Brenton Hordacre, Richard Harris, Julia Pitcher, John Semmler, Nigel C Rogasch, and Mitchell Goldsworthy. Do age-related differences in aperiodic neural activity explain differences in resting eeg alpha? *Neurobiology of Aging*, 121:78–87, 2023.
- Michal Mikl, Radek Marecek, Petr Hlustík, Martina Pavlicová, Ales Drastich, Pavel Chlebus, Milan Brázdil, and Petr Krupa. Effects of spatial smoothing on fMRI group inferences. *Magn. Reson. Imaging*, 26(4):490–503, May 2008.
- Steen Moeller, Pramod Kumar Pisharady, Sudhir Ramanna, Christophe Lenglet, Xiaoping Wu, Logan Dowdle, Essa Yacoub, Kamil Uğurbil, and Mehmet Akçakaya. NOise reduction with DIstribution corrected (NORDIC) PCA in dMRI with complex-valued parameter-free locally low-rank processing. *Neuroimage*, 226:117539, February 2021.
- P. L. Nunez and R. Srinivasan. *Electric Fields of the Brain: The Neurophysics of EEG*. Oxford University Press, New York, NY, 2006.
- Linden Parkes, Ben Fulcher, Murat Yücel, and Alex Fornito. An evaluation of the efficacy, reliability, and sensitivity of motion correction strategies for resting-state functional MRI. *Neuroimage*, 171:415–436, May 2018.
- Judea Pearl. *Causality: models, reasoning and inference*, volume 29. Springer, 2000.
- Joseph Ramsey, Madelyn Glymour, Ruben Sanchez-Romero, and Clark Glymour. A million variables and more: the fast greedy equivalence search algorithm for learning high-dimensional graphical causal models, with an application to functional magnetic resonance images. *Int J Data Sci Anal*, 3(2):121–129, March 2017.
- Joseph D Ramsey, Ruben Sanchez-Romero, and Clark Glymour. Non-Gaussian methods and high-pass filters in the estimation of effective connections. *Neuroimage*, 84:986–1006, January 2014.
- Eric Rawls, Erich Kummerfeld, Bryon A Mueller, Sisi Ma, and Anna Zilverstand. The resting-state causal human connectome is characterized by hub connectivity of executive and attentional networks. *Neuroimage*, 255:119211, July 2022.
- Andrew T Reid, Drew B Headley, Ravi D Mill, Ruben Sanchez-Romero, Lucina Q Uddin, Daniele Marinazzo, Daniel J Lurie, Pedro A Valdés-Sosa, Stephen José Hanson, Bharat B Biswal, Vince Calhoun, Russell A Poldrack, and Michael W Cole. Advancing functional connectivity research from association to causation. *Nat. Neurosci.*, 22(11):1751–1760, November 2019.
- Thomas Richardson. A discovery algorithm for directed cyclic graphs. In *Proceedings of the Twelfth international conference on Uncertainty in artificial intelligence*, UAI’96, pages 454–461, San Francisco, CA, USA, August 1996. Morgan Kaufmann Publishers Inc.
- Ruben Sanchez-Romero and Michael W Cole. Combining multiple functional connectivity methods to improve causal inferences. *J. Cogn. Neurosci.*, 33(2):180–194, February 2021.
- Ruben Sanchez-Romero, Joseph D Ramsey, Kun Zhang, Madelyn R K Glymour, Biwei Huang, and Clark Glymour. Estimating feedforward and feedback effective connections from fMRI time series: Assessments of statistical methods. *Netw Neurosci*, 3(2):274–306, February 2019.
- Alexander Schaefer, Ru Kong, Evan M Gordon, Timothy O Laumann, Xi-Nian Zuo, Avram J Holmes, Simon B Eickhoff, and B T Thomas Yeo. Local-Global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cereb. Cortex*, 28(9):3095–3114, September 2018.
- Shohei Shimizu, Takanori Inazumi, Yasuhiro Sogawa, Aapo Hyvarinen, Yoshinobu Kawahara, Takashi Washio, Patrik O Hoyer, Kenneth Bollen, and Patrik Hoyer. Directing: A direct method for learning a linear non-Gaussian structural equation model. *Journal of Machine Learning Research-JMLR*, 12(Apr):1225–1248, 2011.

Stephen M Smith, Karla L Miller, Gholamreza Salimi-Khorshidi, Matthew Webster, Christian F Beckmann, Thomas E Nichols, Joseph D Ramsey, and Mark W Woolrich. Network modelling methods for FMRI. *Neuroimage*, 54(2):875–891, January 2011.

Kseniya Solovyeva, David Danks, Mohammadsajad Abavisani, and Sergey Plis. Causal learning through deliberate undersampling. March 2023.

Peter Spirtes, Clark N Glymour, Richard Scheines, David Heckerman, Christopher Meek, Gregory Cooper, and Thomas Richardson. *Causation, Prediction, and Search*. MIT Press, 2000.

Martijn P van den Heuvel and Olaf Sporns. Network hubs in the human brain. *Trends Cogn. Sci.*, 17(12):683–696, December 2013.

Luca Vizioli, Steen Moeller, Logan Dowdle, Mehmet Akçakaya, Federico De Martino, Essa Yacoub, and Kamil Uğurbil. Lowering the thermal noise barrier in functional brain mapping with magnetic resonance imaging. *Nat. Commun.*, 12(1):5181, August 2021.

Nora D Volkow, George F Koob, Robert T Croyle, Diana W Bianchi, Joshua A Gordon, Walter J Koroshetz, Eliseo J Pérez-Stable, William T Riley, Michele H Bloch, Kevin Conway, Bethany G Deeds, Gayathri J Dowling, Steven Grant, Katia D Howlett, John A Matochik, Glen D Morgan, Margaret M Murray, Antonio Noronha, Catherine Y Spong, Eric M Wargo, Kenneth R Warren, and Susan R B Weiss. The conception of the ABCD study: From substance use to a broad NIH collaboration. *Dev. Cogn. Neurosci.*, 32:4–7, August 2018.

D J Watts and S H Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, June 1998.

Luc Edward Wilson, Jason da Silva Castanheira, and Sylvain Baillet. Time-resolved parameterization of aperiodic and periodic brain activity. *Elife*, 11:e77348, 2022.