The Heritability of Human Connectomes: a Causal Modeling Analysis

Jaewon Chung[1,*], Eric W. Bridgeford[1], Michael Powell[2], Derek Pisner[3], Ting Xu[4], Joshua T. Vogelstein[1]

Abstract. The heritability of human connectomes is crucial for understanding the influence of genetic and environmental factors on variability in connectomes, and their implications for behavior and disease. However, current methods for studying heritability assume an associational rather than a causal effect, or rely on strong distributional assumptions that may not be appropriate for complex, high-dimensional connectomes. To address these limitations, we propose two solutions: first, we formalize heritability as a problem in causal inference, and identify measured covariates to control for unmeasured confounding, allowing us to make causal claims. Second, we leverage statistical models that capture the underlying structure and dependence within connectomes, enabling us to define different notions of connectome heritability by removing common structures such as scaling of edge weights between connectomes. We then develop a non-parametric test to detect whether causal heritability exists after taking principled steps to adjust for these commonalities, and apply it to diffusion connectomes estimated from the Human Connectome Project. Our findings reveal that heritability can still be detected even after adjusting for potential confounding like neuroanatomy, age, and sex. However, once we address for rescaling between connectomes, our causal tests are no longer significant. These results suggest that previous conclusions on connectome heritability may be driven by rescaling factors. Together, our manuscript highlights the importance for future works to continue to develop data-driven heritability models which faithfully reflect potential confounders and network structure.

1 Introduction

Many common human traits and diseases cluster in families, and are believed to be influenced by several genetic and environmental effects. Heritability quantifies the total variability in a given trait or disease and is attributable to genetic factors and not environmental or stochastic events [1]. Thus, heritability is a key factor in predicting disease risk from an individual's history [2], setting bounds on the ability of genetics to predict disease [2], providing justifications for further genetic studies, and developing genetic treatments for heritable diseases, such as sickle cell anemia [3–7]. There are two main definitions of heritability; narrow-sense heritability (h^2) and broad-sense heritability (H^2) [8]. Narrow-sense heritability measures only the additive genetic effects, and broad-sense heritability measures all possible genetic effects, such as additive, dominance, and epistasis (i.e. the interaction of multiple genes contributing to a phenotype). In typical human twin studies, narrow-sense heritability is most commonly studied because there is insufficient information to estimate the dominant genetic effects and epistasis.

Numerous methods of estimating heritability, that is, partitioning the phenotypic variance into its components (e.g. genetic and environmental), have been applied for various "simple" traits in twin studies. The intraclass correlation (ICC) [9], which estimates the amount of variation that is common to the group as a proportion of the variation in the population, and Falconer's method [8], which is the difference in ICC in identical and fraternal twins of a given trait, have been used to show significant heritability in human personalities [10–12], anatomy [13–16], and intelligence [17, 18]. Structural equation models (SEMs), a type of linear mixed model, have also been used to estimate heritability of phenotypically complex diseases like schizophrenia and Alzheimer's [19, 20]. Some advantages of SEMs are that the model can incorporate multivariate inputs, account for covariates such as sex and age, incorporate familial relationships such as step siblings and parents, and provide more accurate estimates of heritability [21]. Although SEMs are extremely useful in human genetics [22], they exhibit

[1]Johns Hopkins University (JHU)
[2]United States Military Academy (USMA)
[3]Independent Researcher
[4]Child Mind Institute
* corresponding author: j1c@jhu.edu.

several limitations. For example, the phenotypic data are typically assumed to be normally distributed, which is often violated in complex traits, and the phenotypic variance can only be modeled as a linear relationship among its variance components.

With recent advances in neuroimaging techniques, scientists have begun to investigate the heritability of more complex neural traits, such as brain connectivity [23–27]. Connectomes, which are maps of the functional and/or structural connections between distinct brain regions, have proven particularly informative in that they allow for modeling the brain's connectivity topology as a network from the lens of graph theory[28–30]. Bohlken et al. [31], for example, employed SEM using simple graph metrics to study the heritability of white-matter brain network topology in whole-brain structural connectomes, which were derived using diffusion magnetic resonance imaging (dMRI) data. The authors showed both Average Path Length (APL) and Clustering Coefficient (CC) were substantially heritable. Reineberg et al. [32] similarly used SEM to study the heritability of cortical network co-activation patterns in whole-brain functional connectomes derived from functional MRI (fMRI) data. The authors observed that the brain's genetic functional network organization is diverse, whereby <5% of functional links were significantly heritable across both datasets. These studies also exhibited several limitations, such as limited interpretability. It remains unclear, for instance, how to interpret a finding such as greater heritability of APL relative to that of CC. Another limitation alludes to open statistical questions at the intersection of graph theory and connectomics – given that a wide variety graph metrics are available for studying connectomes, how does one select the most appropriate metrics to accurately assess the heritability of brain connectivity, and do so without inflating the risk of false discoveries due to multiple comparisons? Although graph theoretic features are often intuitive and logical descriptors of individual connectomes, no set of graph statistics can comprehensively describe the network topology [33]. Finally, connectomes are inherently non-Euclidean and non-Gaussian data, and current methods, such as SEMs that assume Euclidean and Gaussian data, are inadequate and cannot be applied directly to connectomes. Violating model assumptions can lead to invalid inferences and failure to reproduce.

In this work, we investigate the heritability of human structural connectomes by framing heritability as a causal question; do changes in genetics cause changes in connectomes? We first present and define heritability in a causal model, and present a method to detect and estimate heritability based on distance correlations, which measures linear and non-linear dependence between genetics and connectomes (Figure 1). We then define three models for comparing connectomes based on the statistical modeling of connectomes, which considers the inherent structure and dependencies within networks. Moreover, by explicitly stating the model assumptions, we reveal several possible notions of connectome heritability. For example, we remove common structures with increasing complexity across connectomes and test whether heritability exists beyond these commonalities. Using data from the Human Connectome Project, we demonstrate that structural connectomes are indeed heritable, even after accounting for observed confounding (e.g. age, neuroanatomy, sex). Following the proposed methodology, our findings provide much needed groundwork for future investigations studying the effects of distal genetic influences on complex neurophenotypic traits. Furthermore, our study underscores the importance of considering the inherent structure and dependencies within networks when assessing heritability, as well as the need for methods that can account for non-Euclidean and non-Gaussian data.

2 Methods

For further details on methodological considerations, see Section 5.

2.1 Causal Analysis for Heritability

Throughout this paper, we use P_X to denote a generic distribution function and f_X to denote a generic probability density or mass function for a random variable X, which we abbreviate P and f for simplicity. We use $P_{X|y}$ to denote the conditional distribution function of the random variable X conditioned on $Y = y$. The corresponding conditional density/mass function is $f(Y = y|X = x)$, which we abbreviate $f(y|x)$. We let X denote the $x \in \mathcal{X}$-valued random variable denoting the exposure, Y denote the $y \in \mathcal{Y}$-valued random variable denoting the outcome, Z denote

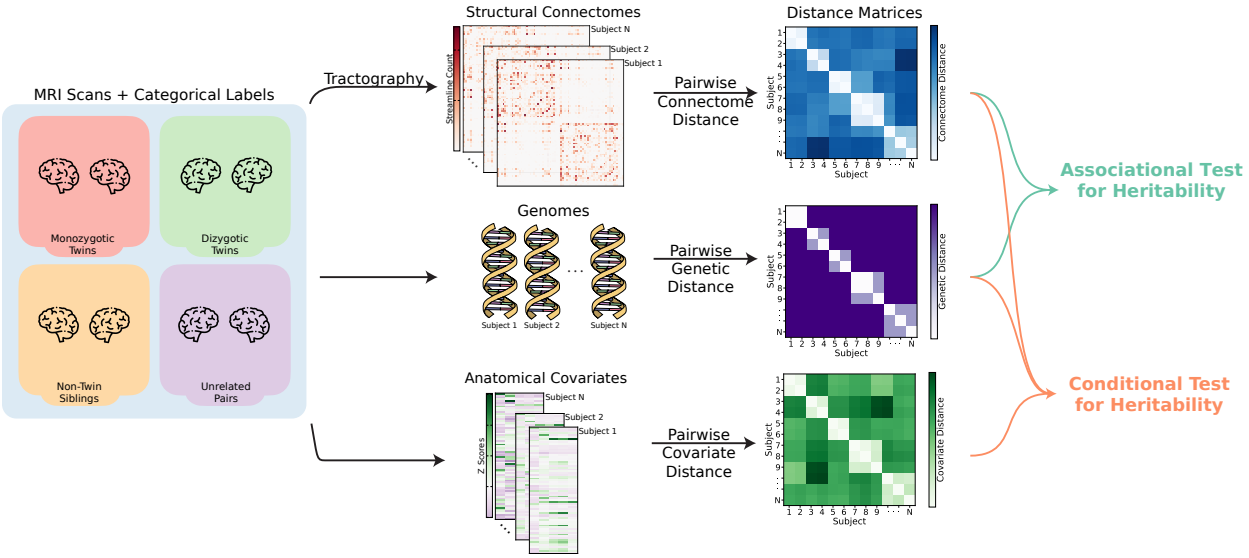

Figure 1: **Overview of the framework for measuring heritability of connectomes.** (*Left*) Diffusion and structural magnetic resonance images (dMRI and sMRI) are processed to generate connectomes. Connectomes are defined on a common set of vertices using brain parcellations. Each pair of individuals has a label corresponding to whether the pair is a monozygotic twin, dizygotic twin, non-twin siblings, or unrelated (e.g. not sharing both mother and father). (*Center left*) For each subject, connectomes and covariates are estimated from MRI scans and phenotypic data. (*Center right*) For each modality (e.g. connectomes, genomes, and covariates), distances are computed between all possible pairs of subjects using the corresponding distance function. This results in distance matrices that are input to subsequent hypothesis tests. (*Right*) Associational effect for connectomic heritability is measured using distance correlation (DCORR). Conditional effect for connectomic heritability is measured using conditional distance correlation (CDCORR). Significance tests provide evidence to reject the null that genomes and connectomes are not related.

the $z \in \mathcal{Z}$-valued random variable denoting the unmeasured covariates, and W denote the $w \in \mathcal{W}$-valued random variable denoting the measured covariates. Throughout, we assume that (X, Y, Z, W) are sampled identically and independently from some true but unknown distribution (which implies exchangeability). In our investigation, the exposure X is the genome conferred to an individual, the outcome Y is an observed trait, such as the structural connectome, and W are covariates regarding the individual which are known (sex, neuroanatomy, and age). For more details on definitions, please see Bridgeford et al. [34].

We are interested in estimating the effect of different exposures on the outcome of interest, which can be quantified using the backdoor formula under the assumption that W and Z close all backdoor paths

$$f_{w,z}(y|x) = \int_{\mathcal{W}\times\mathcal{Z}} f(y|x, w, z) f(w, z) \mathrm{d}(w, z) \tag{2.1}$$

Note that the above equation integrates over the *unconditional* distribution of W and Z, $f(w, z)$. This quantity can therefore be thought of as the distribution of the true outcome averaged over *all* measured and unmeasured covariates. This contrasts from the notationally similar, but practically distinct, quantity $f(y|x)$, e.g.:

$$f(y|x) = \int_{\mathcal{W}\times\mathcal{Z}} f(y|x, w, z) f(w, z|x) \mathrm{d}(w, z).$$

The key difference is that unlike Equation (2.1), $f(y|x)$ averages the true outcome distribution, $f(y|x, w, z)$, over the *conditional* distribution of the measured and unmeasured covariates, $f(w, z|x)$. This has the implication that if $f(w, z|x)$ is not the same as $f(w, z|x')$, $f(y|x)$ represents an average over a different covariate distribution than $f(y|x')$.

Idealized Causal Heritability Given the setup above, we say that *causal heritability* of a trait exists if and only if:

$$f_{w,z}(y|x) \neq f_{w,z}(y) \tag{2.2}$$

and we can construct a hypothesis as follows:

$$H_0 : f_{w,z}(y|x) = f_{w,z}(y) \quad \text{vs} \quad H_A : f_{w,z}(y|x) \neq f_{w,z}(y). \tag{2.3}$$

The null hypothesis states that there is no average causal effect; that is, the genome (exposure) has no effect on the connectomes, and this test can be formalized as an independence test. The key limitation of this characterization of heritability is that the estimand of interest $f_{w,z}(y|x)$ requires an integration over the unconditional distribution of the measured and unmeasured covariates, $f(w, z)$, which we will not have in practice due to the fact that the covariates Z are unknown. Therefore, the above hypothesis test is not feasible since it would not be possible to formulate principled estimators of the estimand $f_{w,z}(y|x)$. In the following sections, we establish a series of related effects while establishing the specific conditions under which the existence of causal heritability effect can be inferred from the effect of interest. It is important to note that the validity of any causal assertions is contingent upon the extent to which these underlying assumptions accurately represent the data.

Associational Heritability We only observe the pairs (x_i, y_i) for $i \in [n]$. Therefore we will only be able estimate functions of (X, Y). We say that *associational heritability* of a trait exists if:

$$f(y|x) \neq f(y)$$

and we can construct a hypothesis as follows:

$$H_0 : f(y|x) = f(y) \quad \text{vs} \quad H_A : f(y|x) \neq f(y). \tag{2.4}$$

Under what circumstances can associational heritability be considered equal to causal heritability? A sufficient assumption for the equivalence of these two effects is that the measured and unmeasured covariates (W, Z) are non-confounding, which is the case when (X, Y) is independent of (W, Z). However, this assumption is often not reasonable. For example, if the size of the human brains are associated to the structural connectomes, then the differences in $f(y|x)$ and $f(y)$ could be due to genetics or neuroanatomy, making it challenging to differentiate between a causal and an associational effect. This motivates the following approach.

Conditional Heritability We observe the triples (x_i, y_i, w_i) for $i \in [n]$. Thus, we will only be able to estimate the functions of (X, Y, W). Thus, a *conditional heritability* of a trait exists if:

$$f(y|x, w) \neq f(y|w)$$

and we can construct a hypothesis as follows:

$$H_0 : f(y|x, w) = f(y|w) \quad \text{vs} \quad H_A : f(y|x, w) \neq f(y|w). \tag{2.5}$$

Under what circumstances can conditional heritability be considered equal to causal heritability? A sufficient condition for these two effects to be equivalent is that the strong ignorability assumption is satisfied [35]. Strong ignorability consists of two conditions:

1. The exposure is independent of the potential outcomes, conditioned on the observed and unobserved covariates (W, Z): $X \perp\!\!\!\perp Y|W, Z$.
2. The distributions of the observed and unobserved covariates have sufficient overlap among the exposure. Specifically, for all $(w, z) \in \mathcal{X} \times \mathcal{Z}$, $0 < \mathbb{P}(X = x|W = w, Z = z) < 1$.

The first condition of strong ignorability can be specified in terms of backdoor paths, which are a set of paths between the exposure and the outcome that are not part of the causal pathway. These paths can create a spurious correlation between exposure and outcome if they are not properly controlled for. Therefore, if all backdoor paths between the exposure and outcome are blocked by conditioning on a set of observed covariates, then the exposure is rendered independent of potential outcomes. In other words, the observed covariates are assumed to be sufficient to control for all confounding variables, including unmeasured covariates, that could influence exposure and outcome. Thus, the main limitation of the conditional effect is that there must be no unmeasured confounding, which is an untestable assumption.

In the context of backdoor paths, the covariate overlap condition ensures that there is adequate variation in the covariates to block all the backdoor paths between the exposures and the outcomes. For example, if there is a subset of the population with extreme values of the covariates where no one is given a particular exposure, then it will be impossible to estimate the effect of the exposure in that subgroup.

Heritability in Twin Datasets The aforementioned effects can be generalized to any trait. In this study, we focus on demonstrating the presence of heritability for structural connectomes in a twin neuroimaging study. Figure 2 illustrates a directed acyclic graph (DAG) underlying a twin study in which measurements are obtained from individuals across multiple families. The genome that an individual inherits from his or her parents is the exposure, and the connectomes are the outcomes. We seek to estimate the causal estimand, which is the causal effect of the genome on the connectome. Both measured and unmeasured covariates can introduce confounding biases that affect both the exposure and the outcome. Unmeasured variables that affect the genome, such as race, fall under the category of **Nature**; unmeasured variables specific to life experiences, including both shared and unshared environmental influences, that affect the connectomes are classified as **Nurture**. Measured covariates, such as **Neuroanatomy**, **Sex**, and **Age**, also have an impact on the connectomes.

In twin studies, it is common to assume that the covariate distributions are the same for monozygotic (MZ) and dizygotic (DZ) twins because they share environmental factors by study design. For example, both MZ and DZ twins often grow up in the same household. This assumption is known as the equal environment assumption and is often used to estimate genetic and environmental contributions to a trait or outcome of interest. It is also reasonable to assume that covariate distributions still overlap to some extent in non-twin siblings, especially if they grew up in the same household or had similar life experiences. Given proper experimental design considerations to allow for other covariate distributions, such as age distribution to overlap, we will assume that the covariate set of neuroanatomy, sex, and age close all backdoor paths.

Hypothesis Tests for Discovering Heritability Classical statistical tests such as Student's t-test, one-way analysis of variance (`ANOVA`), and their multivariate counterparts, such as Hotelling's T^2 test and multivariate analysis of variance (`MANOVA`), are not well-suited for testing heritability in twin studies due to the inherent structure of the data. Specifically, the familial relationships (e.g. monozygotic twins, dizygotic twins) depend on which pair of subjects being examined rather than the individuals. For instance, consider two families, each with a set of monozygotic twins. Within families, the subjects are considered as monozygotic twins, but when compared across families, they are considered as unrelated subjects. Consequently, although both pairs are classified as monozygotic twins, they cannot be combined into a group for subsequent analysis. To address this issue, some studies examine within family pair differences rather than individual observations. However, this approach introduces another challenge: the pairwise differences are not independent. For example, in a family with one monozygotic twin pair and a non-twin sibling, there are three possible pairings (one twin pair and two non-twin sibling pairs). When comparing the connectomes of all possible pairs among the three individuals, the twins would be used for comparison in both the monozygotic and sibling groups.

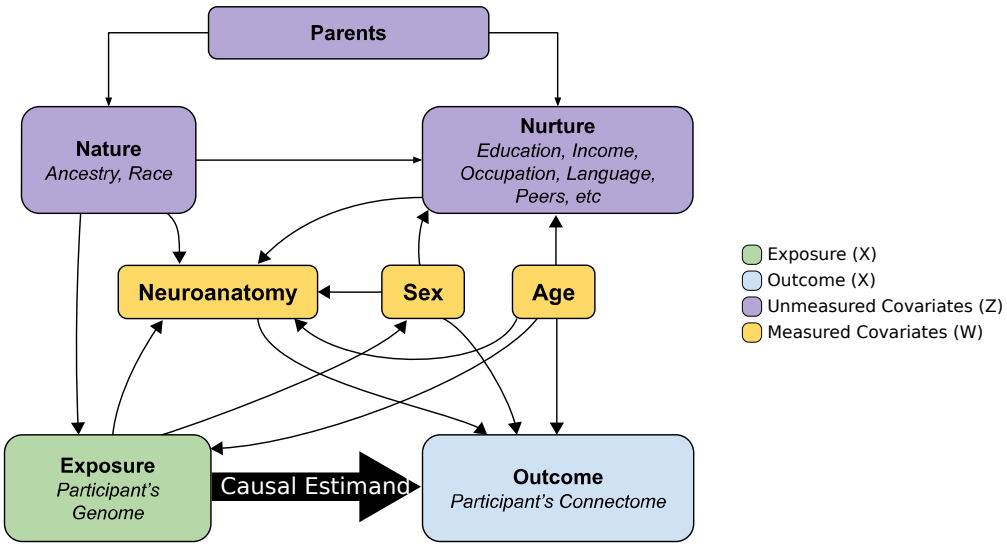

Figure 2: **Causal Graph of Study Covariates.** Causal directed acyclic graph (DAG) presents a visual representation of the potential relationships between the genome and connectome, in the context of a twin study. The descriptions depict the various attributes that may be considered as exposures, outcomes, and covariates in such a study. In the absence of confounding, associational effects can be considered as causal effects. However, conditional effects can only be considered as causal when the covariates measured are sufficient to close all backdoor paths.

Distance correlation (DCORR) is a non-parametric test that can detect linear and non-linear dependency between two multivariate variables, and resolves the issue of pairwise dependence [36, 37]. The first step in DCORR is to compute the distance matrices for the variables being compared. These matrices represent the distances, or differences, between all possible pairs of observations in the dataset. By working with distance matrices, DCORR inherently takes into account the relationships between pairs of observations. Therefore, we used distance correlation to test the hypothesis of associational effect (Eq. 2.4). To test the hypothesis of causal effect (Eq. 2.5), we use the conditional distance correlation (CDCORR), which augments the DCORR procedure by conditioning on the kernel of third variable [38]. In the following section, we present models and methods of computing distances between a pair of connectomes.

2.2 Statistical modeling for connectomes The random dot product graph (RDPG) model provides a way to model weighted networks and compute differences between them using statistically principled procedures [39–42]. In this model, a vertex is a region-of-interest (ROI) in the brain, which is represented as a low-dimensional vector called a latent position. The probability of one ROI connecting to another is determined by the dot product of the corresponding latent positions. In other words, a matrix containing the latent positions of all ROIs is a representation of the underlying distribution of the connectome. Given these representations, we propose three different models for computing distances between connectomes (see Section 5.4 for more details):

1. **Exact model**: This model measures all differences between latent positions, with differences in the latent positions implying differences in the connectomes themselves. Consequently, a hypothesis test based on an exact model aims to test whether there is a correlation between the actual differences in connectomes and the differences in the genomes.
2. **Global model**: This model examines whether the latent positions of one connectome are a scaled version of the other. For example, if the number of edges in male connectomes are consistently larger than those in females, we have no way of differentiating whether significant findings from the exact model are a result of differences in scaling or differences in the fundamental structure of the connectomes themselves. Thus, hypothesis tests based on global model test whether differences in connectomes after adjusting for scaling differences correlate

with differences in genomes. The global model assumes that the scaling factor is the same across all regions, meaning that any differences in scaling are not region-specific.

3. **Vertex model**: This model is similar to the global model, but it allows for each vertex to be scaled differently. The idea behind this approach is that some vertices may have a greater impact on the overall network than others, so scaling them differently can provide a more accurate representation of the network. Consider the examples of how brain regions connect with each other. Regions in the same hemisphere are more likely to be connected than across hemispheres. Even within the same hemisphere, different regions may have distinct preferences for forming connections with other specific regions. Thus, hypothesis tests based on the vertex model aim to determine whether there are significant differences in connectomes after accounting for differences in vertex-wise scaling.

To demonstrate the differences in the connectome models, we simulated connectomes with various latent positions (Figure 3. Columns I and II display four pairs of connectomes, each with specific latent position models as explained in **??**. Rows (A-C) share the same underlying latent positions, while row (D) has different latent positions. When the connectome model can account for the differences in latent positions between the subjects, we observe a "uniform" structure in the differences, with no visible pattern. For the exact model (column III), distances increase as we move down the rows because it cannot account for any scaling, whether global or vertex-wise. We observe patterns in the differences, where red indicates edges in which subject 1 tended to have more edges, while blue shows edges where subject 1 tended to have fewer edges. For instance, for each simulation, when subject 1 tended to have more edges (row B.I-II), we observe uniform red coloring in the differences (row B.III). Similarly, for the global model, we only see patterns in differences when the differences in latent positions are due to vertex scaling, as shown by the gradients of red and blue (C.IV and D.IV). The main takeaway is that any hypothesis tests based on these models can only make claims within the limits of their ability to explain the differences in the pairs of connectomes.

3 Results

3.1 Dataset We used an open access human brain diffusion (dMRI) and structural magnetic resonance imaging (sMRI) data set from the Human Connectome Project (HCP) Young Adult study, acquired by Washington University in St. Louis (WUSTL) and the University of Minnesota (Minn) [43, 44]. Out of the 1206 participants, a total of 1024 participant brain images were processed. Families with only one subject were discarded, and half-siblings were discarded. The data set includes 322 (196 females; ages $22-36$ years old, mean = 29.6 and SD = 3.3) monozygotic twins, 212 (125 females; ages $22-36$ years old, mean = 28.6 and SD = 3.4) dizygotic twins, and 490 (237 females; ages $22-37$ years old, mean = 28.3 and SD = 3.9) non-twin siblings. Structural connectomes were derived based on various parcellations. See 5.1 for more details on network construction. Figure 4 shows visualizations of the average connectome and the difference in two most different subjects for the Desikan parcellation.

3.2 Associational Tests for Connectomic Heritability We begin the examination of heritability by first comparing connectomes from monozygotic and same-sex dizygotic twins. By limiting the comparison to these two groups, we minimized the impact of environmental factors (e.g. Nurture in Figure 2) on the connectomes and assessed the ability of connectomes to capture variations in the genome. Shared and non-shared environmental influences may arise due to differences in age or sex, which can lead to differences in upbringing or educational experiences. By controlling for such factors, we could reasonably attribute any differences in connectomes to variations in genomes rather than environmental factors. Moreover, if the differences in connectomes between monozygotic twins were "smaller" than those between dizygotic twins, it would suggest that connectomes can indeed capture variations in genomes. This would support the idea that variations in the genome can manifest as differences in the connectomes, and thus provide insights into the genetic basis of brain connectivity.

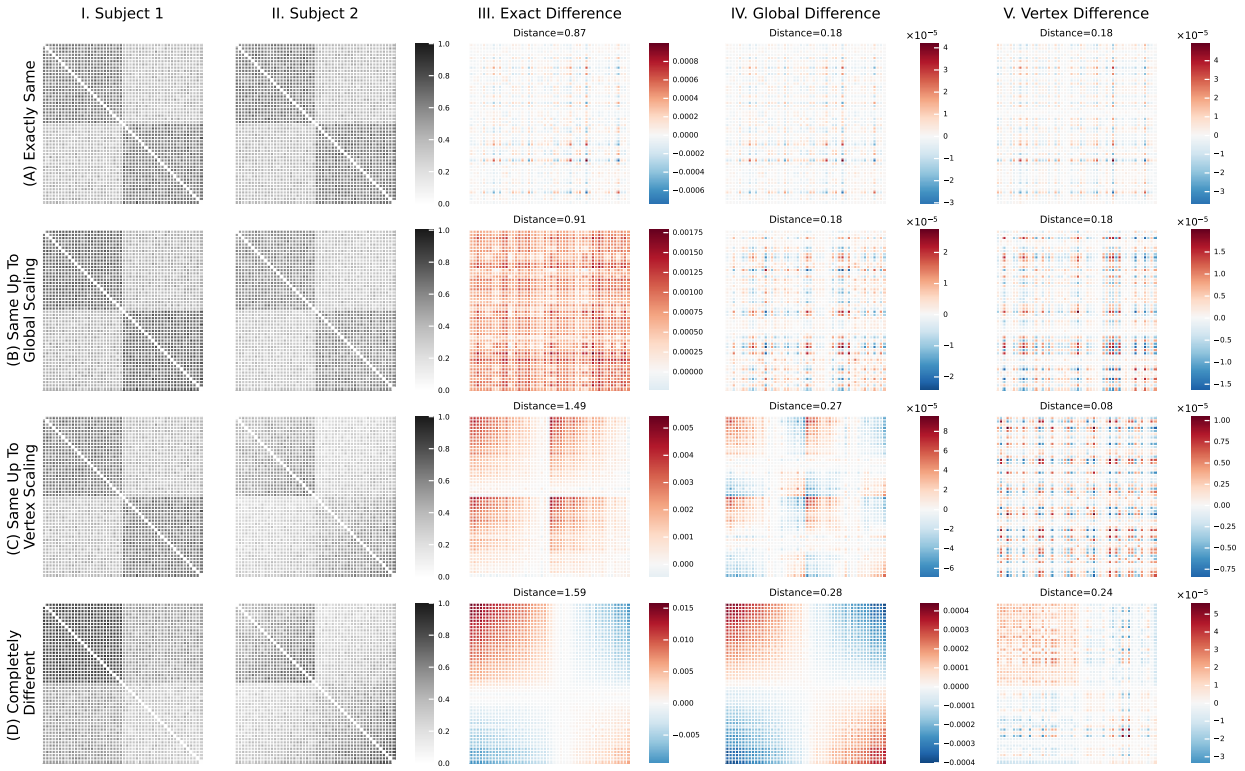

Figure 3: **Simulations demonstrate the utility of connectome models.** Connectomes with 50 vertices were simulated 100 times. Columns (I-II) display the average simulated connectomes, while columns (III-V) illustrate the differences in latent positions averaged over these simulations. In columns (III-V), red corresponds to when subject 1 had more edges than subject 2, while blue corresponds to when subject 1 had fewer edges. Rows (A-D) correspond to various parameters used to generate the connectomes. Connectomes in rows (A-C) share the same latent positions, subject to scaling transformations, whereas connectomes in row (D) have distinct latent positions, conditional on scaling transformations. When the connectome models accurately represent the generative process, the difference matrix displays a "uniform" structure without discernible patterns (A.III-V, B.IV-V, C.V). However, when the connectome models are not appropriate, patterns emerge, as indicated by the increased distance values (B.III, C.III-VI, D.III-V). For further technical details of these simulations, see Appendix **??**.

The distribution of connectome distances is shown in Figure 5A.I-III by the blue and yellow rows, revealing that the medians of monozygotic twins ($N = 127$) are smaller than those of same-sex dizygotic twins ($N = 75$) for all three connectome models. To determine the significance of this difference, we perform a test for an associational effect using DCORR, yielding p-values of $< 1 \times 10^{-3}$ for all models (Figure 5B.I). These results provide our first evidence for the heritability of connectomes. We then expanded our examination to include non-twin siblings and unrelated subjects (Figure 5A.I-III green and red rows) and observed that the medians grow larger as the familial relationships grow distant. Based on these observations, we tested for associational heritability across all familial relationships, yielding p-values of $< 1 \times 10^{-3}$ for males and females under all three heritability models (Figure 5B.II). All of these tests indicate a strong associational effect of the genome on connectomes.

3.3 Associational Test for Neuroanatomic Heritability

Multiple studies have investigated the heritability of human brain anatomy [45–50]. Structural MRI studies on adults have shown that genetic and environmental factors have a significant impact on brain properties such as volume, shape, and both cortical and subcortical structures. This suggests that neuroanatomy could be an effect mediator (Figure 2). To assess if neuroanatomy could be a mediator, several features such as brain volume, axial diffusivity (AD), radial diffusivity (RD), and fractional anisotropy (FA) were computed for each region of the brain. Refer to Section 5.4 for more details.

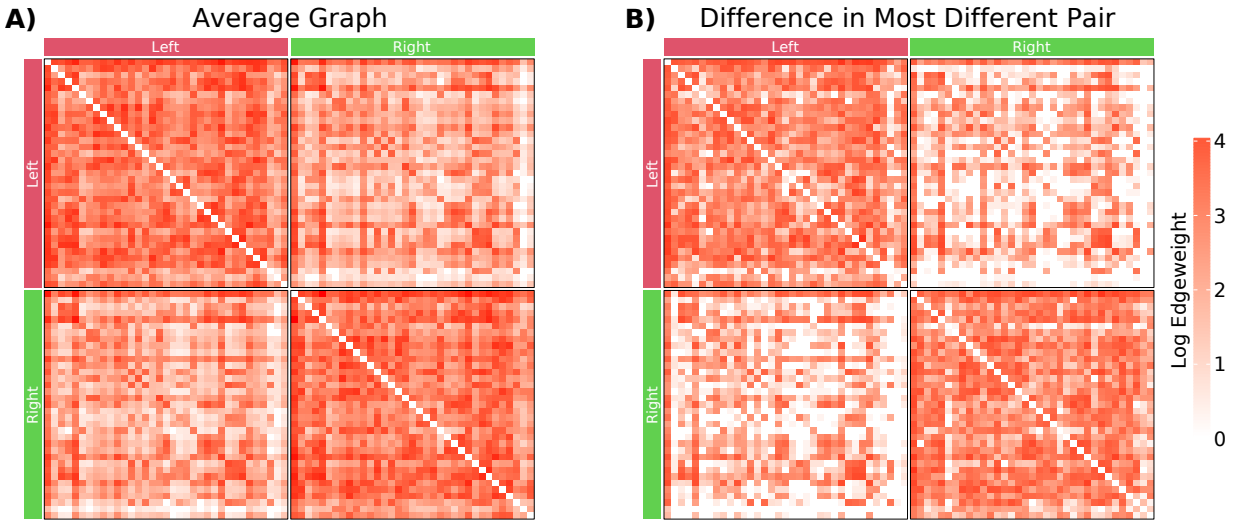

Figure 4: **Visualization of connectomes as adjacency matrices using the projected Desikan parcellation with hemispheric labels.** The projected Deskian parcellation was chosen for visualization purposes due to the low number of vertices ($n = 70$). **(A)** Average connectome of all subjects with log-transformed edge weights. **(B)** Absolute difference of connectomes from the most different pair of subjects with log-transformed edge weights.

We used DCORR to test for the associational effect of the genome on anatomical covariates. This test assumes that if the covariates and genetics are independent, then brain anatomy is not heritable and has no effect on connectomes. Conversely, if the alternative is true, the associational effects of heritability of connectomes may be partially explained by the heritability of the anatomy itself, making the brain anatomy an effect mediator. Figure 5A.IV illustrates that similar to connectome distances, the median distances increase as familial relationships become more distant. Furthermore, we find that the tests yielded p-values of $< 1 \times 10^{-3}$ for all three gender groups under all three connectome models (Figure 5B.III). The significance suggest that the associational heritability of connectomes was partly due to the effect of anatomy on connectomes.

3.4 Causal test for connectomic heritability To account for the dependence of neuroanatomy on genetics, we utilized conditional distance correlation (CDCORR) to estimate the conditional effect of genome on connectomes (Eq. 2.5). The covariate set also included age, which is a confounder. A non-significant result implies that similarities in brain connectivity are completely explained by similarities in brain anatomy and age, whereas a significant result suggests that the connectomes contain information on brain connectivity beyond anatomy and age, which is still heritable. Our results showed significant p-values of $< 10^{-2}$ for exact and global models, but non-significant p-values for the vertex model. This indicates that while connectomes are heritable and not entirely influenced by anatomy and age, there is no significant underlying structure beyond the vertex-wise scaling. This contrasts with the results shown in Figure 5B.I-II, highlighting the importance of causal modeling in this study.

3.5 Non-heritable Subgraphs are also Heritable In this section, we consider the possibility that heritability effects detected in Section 3.4 are driven by a subset of vertices rather than the entire connectome. That is, it is possible that the connectivity of a subset of brain regions may be highly similar in twins but highly different in non-twin siblings or unrelated individuals. To investigate this, we test for causal heritability by examining the *non-heritable subgraphs*, which are subgraphs induced by the vertices in which no heritable effect can be detected. To identify the non-heritable subgraphs, we modify the hypothesis in Eq. 2.5 to allow for vertex-wise testing:

$$H_0 : f(y^k|x^k, w^k) = f(y^k|w^k) \quad \text{vs} \quad H_A : f(y^k|x^k, w^k) \neq f(y^k|w^k). \tag{3.1}$$

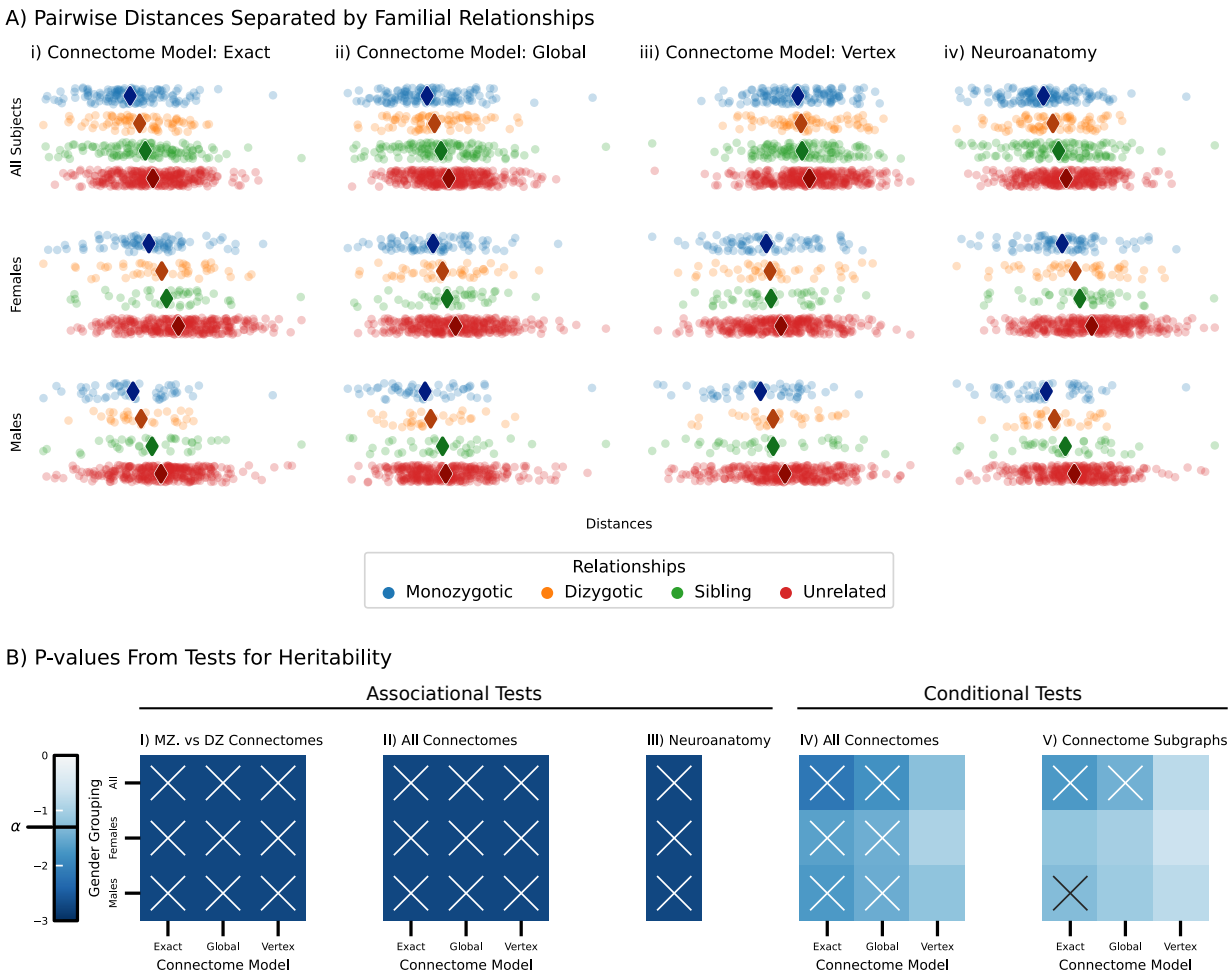

Figure 5: **Testing associational and causal effect of genome on connectomes and neuroanatomy. (A)** Each point represents the pairwise distance between pairs of participants; diamond markers represent the median distance, colors are familial relationships, and rows are sex. **(A.i-iii)** The median distances revealed that MZ twins exhibited smaller distances compared to DZ twins, siblings, and unrelated pairs, whereas the unrelated pairs showed larger distances. DZ twins displayed smaller distances compared to siblings and unrelated pairs, but larger than those of MZ twins. These findings provide qualitative evidence of heritability of connectomes. **(A.iv)** The median distances of neuroanatomy also show qualitative evidence of heritability. **(B)** Colors of heatmaps represent p-values, rows are gender groups, and columns are connectome models. (B.I-II) Associational test for connectome heritability. In (B.I), we only compare MZ and DZ twins, and the significance tests provide evidence of heritability as shown in (A.I-III). In (B.II), we conduct tests across all three familial relationships and observe significant effects, providing further evidence for connectomic heritability. In contrast, in (B.III), we observe significant effects for neuroanatomy, suggesting that the results in (B.I-II) may be confounded by an effect mediator. Subsequently, in (B.IV-V), we perform causal tests for connectome heritability. In (B.IV), we observe significant results only for the exact and global models, indicating that the connectomes remain heritable even after conditioning on the anatomical covariates, age, and sex. This suggests that the results in (B.I-II) cannot be entirely explained by neuroanatomy. The non-significant tests using the vertex model suggest that when we control for covariates and vertex-wise scaling, we can no longer detect heritability. In (B.V), we define a non-heritable subgraph as a collection of vertices that are not heritable by the causal test. We still observe several significant tests, suggesting heritability remains even after discarding vertices.

where k indexes a particular vertex, and y^k is a matrix of latent positions from all connectomes for vertex k. After computing a p-value for each vertex, those with p-values less than $\alpha = 0.05$ after multiple hypothesis correction were discarded. The remaining set of vertices induces the non-heritable subgraphs. We then tested for conditional heritability of the non-heritable subgraphs as in Eq. 2.5.

The outcomes presented in Figure 5B.V display significant p-values for both the exact and global models across all genders, with only the exact model presenting significant p-values in males. These

results suggest that for females, heritability can be fully accounted for by the heritable vertices, while in males, the non-heritable vertices still demonstrate heritability. Nevertheless, when controlling for scaling, the significance observed in the exact model for males is eliminated, indicating that differences in the scaling of the subgraphs explain the observed significance. These findings suggest the presence of an underlying structure in the non-heritable subgraph that cannot be detected at the level of individual vertices and highlights the necessity of statistical modeling of networks.

4 Discussion

4.1 Summary In this study, we presented the first examination of the heritability of human connectomes using causal modeling and statistical modeling of connectomes. Our investigations into the heritability of connectomes provides strong evidence that the structural connectivity patterns within them are highly influenced by genetics by leveraging causal models and developing a test procedure to detect heritability. To establish this, we first compared the connectomes of monozygotic twins with same-sex dizygotic twins, which allowed us to account for shared environmental factors. Our results demonstrated that not only are connectomes reliable and meaningful, but they are also significantly influenced by genetics. We then extended our analysis to include monozygotic, dizygotic, and non-twin siblings, which further confirmed that connectomes are heritable. However, these results may be confounded by both observed and unobserved covariates. Therefore, we define a set of covariates, which we believe can account for all possible other unobserved covariates. Given this, we found that connectomes remained heritable even after controlling for age and the mediating effect of neuroanatomy. Finally, we demonstrated that the subgraphs formed by sets of brain regions that are not heritable within connectomes are also heritable. This suggests that there is an underlying structure within these subgraphs that is not evident in individual vertices.

A key innovation in this study is the utilization of statistical models, specifically the random dot product graphs, for connectomes, which introduces a new representation called latent positions. The models allow us to compare connectomes in various ways and explore different aspects of heritability. The simplest model we use is the exact model, which assesses any differences in the latent positions of connectomes. The global model, on the other hand, accounts for the possibility that connectomes may have different scales, such as variations in edge weights or more connections in general. Therefore, this model quantifies the differences in latent positions after removing the effect of global scaling. The vertex model allows for both global and vertex-wise scaling, enabling the three models to account for increasingly complex structures in connectomes as we move from the exact model to the vertex model. Our findings, as presented in Section 3.4, indicate that genetics play a crucial role in the differences of vertex scaling. However, when addressing vertex scaling of connectomes, heritability becomes undetectable, suggesting that heritability is only due to vertex scaling and not other underlying structures present in the connectomes. These results imply that previous conclusions on connectomic heritability may be influenced by vertex scaling and emphasize the need for further research into the relationship between vertex scaling in RDPG and network statistics, such as clustering coefficients and average path lengths.

There are several other open questions and opportunities for future research on the heritability of connectomes. One potential direction involves studying heritability in different neuroimaging modalities. Although our focus was on human connectomes estimated from diffusion MRI, the methods presented here can be applied to study heritability in connectomes obtained from other modalities, such as functional MRI, magnetoencephalography (MEG), and electroencephalography (EEG), or even in other species, like mice. Another potential direction is incorporating more detailed genetic data, including GWAS, to identify specific genetic variants associated with brain connectivity patterns. This would enable more direct and meaningful comparisons of genomes between subjects and could be readily integrated into the testing procedure presented here.

The findings on heritability in connectome analysis have significant implications for understanding

the genetic basis of various neurological and psychiatric disorders. Disruptions in brain connectivity have been implicated in numerous conditions, such as autism spectrum disorder, schizophrenia, and Alzheimer's disease [51, 52]. By revealing the heritable components of brain connectivity, researchers can start to pinpoint specific genetic factors that contribute to the development of these disorders. This knowledge could inform the creation of novel diagnostic tools, treatment strategies, and targeted interventions, ultimately enhancing the prognosis and quality of life for individuals affected by these disorders.

4.2 Limitations

Estimating Heritability Heritability is typically defined as the relative contribution of genetic variations as opposed to environmental variations in phenotypic traits, and tests whether the heritability estimate is statistically significant, meaning if heritability exists or not. Numerous methods have been developed to partition phenotypic variation into genetic and environmental components, including techniques such as structural equation modeling (SEM). However, the causal approach presented in this study is distinct from these methods, as it does not aim to estimate heritability in a quantitative manner. Instead, our approach focuses on identifying the presence or absence of heritability by examining causal relationships between genetic factors and phenotypic traits.

The causal approach presented in this study offers a complementary perspective that can contribute to resolving the missing heritability problem. The missing heritability problem arises when the estimated heritability from genetic markers, such as those identified in genome-wide association studies (GWAS), accounts for only a small fraction of the heritability inferred from twin studies using SEMs or similar methods. This discrepancy between the observed and expected heritability is a major challenge in genetic research, as it limits our understanding of the genetic architecture underlying complex traits and diseases. By focusing on the presence or absence of causal relationships between genetic factors and phenotypic traits, this approach can help to identify previously unexplored genetic influences that might not be captured by traditional heritability estimation methods.

Choice of Model We utilized the random dot product graph (`RDPG`) model in our analysis, which provided a unified framework for analyzing connectomes at both the network and individual vertex levels. However, it is important to acknowledge that every model has its limitations. The `RDPG` model can only represent a subset of all possible stochastic block models, and may not be able to capture certain patterns in connectivity. For example, if human brains were characterized by more connections between hemispheres than within hemispheres, the `RDPG` model will not be able to capture these patterns. To address these limitations, other models have been developed, such as the generalized random dot product graph, which can represent a broader range of stochastic block models and other network models. Additionally, there are models that are designed to analyze populations of networks, such as the common subspace independent edge model and joint random dot product graph model. Choosing any of these models enables one to define a valid distance metric for comparing connectomes. It is worth noting that there is currently no consensus on which model is best suited for modeling the human brain. Therefore, researchers must carefully consider the advantages and limitations of each model when selecting the appropriate one for their analysis. Ultimately, the choice of model will depend on the research question, the data available, and the assumptions that are deemed reasonable for the specific context.

5 Methods

5.1 Network Construction

Connectomes were estimated using the `ndmg` pipeline [53], which is designed to process and generate connectomes from human dMRI and sMRI that minimizes batch effects across datasets. We note that higher b-values than 1000 were discarded prior to preprocessing. The dMRI scans were first denoised from Gibbs ringing using DiPy's `gibbs-removal`, and then were pre-processed for eddy currents using FSL's `eddy-correct` [54]. The results were denoised using

DiPy's `Local PCA`. FSL's "standard" linear registration pipeline was used to register the sMRI and dMRI images to the MNI152 atlas [54–57]. A constant solid angle orientation distribution function (CSA-ODF) model is fit using DiPy [58] to obtain an estimated tensor at each voxel. A local deterministic tractography algorithm is applied using DiPy's [58] to obtain streamlines, which indicate the voxels connected by an axonal fiber tract. Given a parcellation with vertices V and a corresponding mapping $P(v_i)$ indicating the voxels within a region i, we contract our fiber streamlines as follows. $w(v_i, v_j) = \sum_{u \in P(v_i)} \sum_{w \in P(v_j)} \mathbb{I}\{F_{u,w}\}$ where $F_{u,w}$ is true if a fiber tract exists between voxels u and w, and false if there is no fiber tract between voxels u and w. Several parcellations were used to generate connectomes, such as the Schaefer [59], AAL [60], Glasser [61] and a projection of the Desikan parcellation [62] to sub-cortical white-matter structures [63].

5.2 Network Preliminaries Networks (or graphs) are convenient mathematical objects for representing connectomes. A network G consists of an ordered set (V, E) where V is the vertex set, and E is the set of edges. The set of vertices is represented as $V = \{1, 2, \ldots, n\}$ where $|V| = n$. The set of edges is a subset of all possible edges (e.g. $E \subset V \times V$). An edge exists between vertex i and j if $(i, j) \in E$. A network can also be represented by its adjacency matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ where each element, $\mathbf{A}_{ij}$, represents the edge between vertices i and j. In many connectomics datasets, edges have associated weights that represents some notion of strength of connections between two vertices. For example, edge weights from structural connectomes are non-negative integers that represent the number of estimated fiber tracts between two brain regions.

5.3 Statistical Model for Connectomes Connectomes can be modeled using statistical models designed for network data. These statistical models consider the entire network as a random variable, including the inherent structure, dependencies within networks, and the noise in observed data [33, 64, 65]. In the following sections, we define the *random dot product graph* (`RDPG`) [40, 66], how the parameters can be estimated from observed data, and then define the testing procedure.

Random Dot Product Graph Model The Random Dot Product Graph (`RDPG`) is a type of independent edge model. In this model, each element of an adjacency matrix is sampled independently from a Bernoulli distribution:

$$\mathbf{A}_{ij} \sim \text{Bernoulli}(\mathbf{P}_{ij})$$

Given the number of vertices n, the matrix $\mathbf{P}$ is a $n \times n$ matrix of edge-wise connection probabilities with elements in $[0, 1]$. We can construct various models depending on the constraints imposed on $\mathbf{P}$. Note that we assume that $\mathbf{P}$ has no self-loops (i.e. $\text{diag}(\mathbf{P}) = \vec{0}$) and is undirected (i.e. $\mathbf{P}^\top = \mathbf{P}$).

In the random dot product graph (`RDPG`), the probability of a connection $\mathbf{P}_{ij}$ is determined by the vertices. Each vertex $i \in V$ is associated with a low-dimensional *latent position* vector, $\mathbf{X}_i$, in the Euclidean space $\mathbb{R}^d$. The probability of connection between vertices i and j is given by the dot product (i.e $\mathbf{P}_{ij} = \mathbf{X}_i \mathbf{X}_j^\top$). Thus, in a d – dimensional `RDPG` with n vertices, the rows of the matrix $\mathbf{X} \in \mathbb{R}^{n \times n}$ are the latent positions of each vertex, and the matrix of edge-wise connection probabilities is given by $\mathbf{P} = \mathbf{X}\,\mathbf{X}^\top$. Each element of the adjacency matrix is then independently modelled as

$$\mathbf{A}_{ij} = \text{Bernoulli}(\mathbf{X}_i \mathbf{X}_j^\top)$$

where $\mathbf{X}_i$ and $\mathbf{X}_j$ are latent positions for vertices i and j, respectively.

We acknowledge that the original intention of `RDPG` is to model binary networks, although the model can be naturally extended to handle weighted networks. However, the weighted `RDPG` models are not well studied, and as such does not enjoy the same statistical guarantees. In the subsequent section, we present an algorithm for estimating latent positions from observed data and describe methods for preprocessing the data to enable interpretation of results within the context of binary networks while still utilizing weighted network data.

Adjacency Spectral Embedding The modeling assumptions of the RDPG make the estimation of latent positions, which are usually unobserved in practice, analytically tractable. The estimation procedure we use is *adjacency spectral embedding* (ASE) [67]. The ASE of an adjacency matrix $\mathbf{A}$ into d-dimensions is $\widehat{\mathbf{X}} = \widehat{\mathbf{U}}|\widehat{\mathbf{S}}|^{1/2}$, where $|\widehat{\mathbf{S}}|$ is a diagonal $d \times d$ matrix containing on the main diagonal the absolute value of the top-d eigenvalues of $\mathbf{A}$ in magnitude, in decreasing order, and $\widehat{\mathbf{U}}$ is an $n \times d$ matrix containing the corresponding orthonormal eigenvectors. This simple and computationally efficient approach results in consistent estimates $\widehat{\mathbf{X}}$ of the true latent positions $\mathbf{X}$ [67–69]. The ASE depends on a parameter d that corresponds to the rank of the expected adjacency matrix conditional on the latent positions; in practice, we estimate this dimension, $\hat{d}$, via the scree plot of the eigenvalues of the adjacency matrix which can be done automatically via a likelihood profile approach [70].

Structural connectomes have weighted edges (i.e. non-negative integers) and have no self-loops. To improve estimation of latent positions via ASE, we implement two preprocessing steps in order: 1) *pass-to-ranks* (PTR), and 2) *diagonal augmentation* (DA). PTR is a method for rescaling the *positive* edge weights such that all edge weights are in $[0, 1]$. The motivation of PTR is to convert the edge weights such that its distribution is uniform [53, 71]. Given an adjacency matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, let $R(\mathbf{A}_{ij})$ be the "rank" of $\mathbf{A}_{ij}$, that is, $R(\mathbf{A}_{ij}) = k$ if $\mathbf{A}_{ij}$ is the kth smallest element in $\mathbf{A}$. The rescaled adjacency matrix, $\widetilde{\mathbf{A}}$, is:

$$\widetilde{\mathbf{A}}_{ij} = \begin{cases} \frac{R(\mathbf{A}_{ij})}{|E|} & \text{if } \mathbf{A}_{ij} > 0 \\ 0 & \text{otherwise} \end{cases}$$

where $|E|$ is the number of non-zero edges. Ties in rank are broken by averaging the ranks. DA is a method for augmenting the diagonal of an adjacency matrix. Given an adjacency matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, the diagonal augmented matrix $\bar{\mathbf{A}} = \mathbf{A} + \mathbf{D}$ where $\mathbf{D} \in \mathbb{R}^{n \times n}$ is a diagonal matrix with entries $\mathbf{D}_{ii} = \sum_{j=0}^{n} \mathbf{A}_{ij}$.

5.4 Distance Measures for Distance Correlation and Conditional Distance Correlation

The hypothesis testing procedures in Section 2.1 require functions to quantify dissimilarity (distances) or similarities between pairs of observations. In the following sections, we detail the functions used to compute the distance matrices for connectomes and genomes, and the kernel matrix for covariates.

Distances of Connectomes In [41], the authors proposed three methods for calculating the distance between two networks under the RDPG model. Rather than comparing the high-dimensional connectomes directly, we first obtain low-dimensional latent positions of the connectomes via ASE and then compared these new representations. Each method makes different assumptions about the structure of the connectomes, providing different interpretations of the results.

Let $\mathbf{A}^{(1)}$ and $\mathbf{A}^{(2)}$ be $n \times n$ adjacency matrices, and let $\widehat{\mathbf{X}}$ and $\widehat{\mathbf{Y}}$ be $n \times d$ latent position matrices from their respective ASE. The distance metrics are defined as:

1. **Exact Model**:

$$d(\widehat{\mathbf{X}}, \widehat{\mathbf{Y}}) = \min_{\mathbf{W}} \|\widehat{\mathbf{X}}\,\mathbf{W} - \widehat{\mathbf{Y}}\|_F \tag{5.1}$$

   where $\mathbf{W}$ is a $d \times d$ orthogonal matrix that is estimated by solving the orthogonal Procrustes problem.

2. **Global Model**:

$$d(\widehat{\mathbf{X}}, \widehat{\mathbf{Y}}) = \min_{\mathbf{W}} \frac{\left\| \widehat{\mathbf{X}}\,\mathbf{W} / \left\|\widehat{\mathbf{X}}\right\|_F - \widehat{\mathbf{Y}} / \left\|\widehat{\mathbf{Y}}\right\|_F \right\|_F}{1/\left\|\widehat{\mathbf{X}}\right\|_F + 1/\left\|\widehat{\mathbf{Y}}\right\|_F} \tag{5.2}$$

   Note that a scaling constant is not estimated directly but the effect of scaling is removed by ensuring all latent position matrices have the same scale by normalizing by their respective Frobenius norms.

3. **Vertex Model**: For any matrix $\mathbf{Z}$ with shape $n \times d$, we define the $\mathcal{D}(\mathbf{Z})$ as a diagonal $n \times n$ matrix whose main diagonal entries are the Euclidean norm of the rows of $\mathbf{Z}$. Let $\mathcal{D}^{-1}(\mathbf{Z})$ be defined as $1/(\min_i \|\mathbf{Z}_i\|)$, and let $\mathcal{P}(\mathbf{Z})$ be the matrix whose rows are the projection of the rows of $\mathbf{Z}$ onto the unit sphere. Then the distance is defined as:

$$d(\widehat{\mathbf{X}}, \widehat{\mathbf{Y}}) = \min_{\mathbf{W}} \frac{\left\| \mathcal{P}(\widehat{\mathbf{X}})\, \mathbf{W} - \mathcal{P}(\widehat{\mathbf{Y}}) \right\|_F}{\left\| \mathcal{D}^{-1}(\widehat{\mathbf{X}}) \right\|_2 + \left\| \mathcal{D}^{-1}(\widehat{\mathbf{Y}}) \right\|_2} \tag{5.3}$$

Distances of Genomes Genomes are costly and challenging to measure, but their similarity can be estimated even without measurements by making strong assumptions. The coefficient of kinship (ϕ_{ij}) is commonly used, which represents the probability that two individuals (subjects i and j) have the same allele or a DNA sequence at a random position in the genome [72]. For instance, monozygotic twins have an identical genome but two possible alleles since humans have two copies of a gene, resulting in a coefficient of $\phi_{ij} = 1/2$. On the other hand, dizygotic twins and siblings share 50% of the genome on average, resulting in a coefficient of $\phi_{ij} = 1/4$. Therefore, the distance between two genomes G_i and G_j is defined as:

$$d(G_i, G_j) = 1 - 2\phi_{ij}$$

Table 1 enumerates all the possible distances. In the context of distance correlations, we note that any monotonic transformations of the values given to the genetic distance are equivalent as long as the ordinality remains.

Relationship	ϕ_{ij}	$1 - 2\phi_{ij}$
Monozygotic	$\frac{1}{2}$	0
Dizygotic	$\frac{1}{4}$	$\frac{1}{2}$
Siblings	$\frac{1}{4}$	$\frac{1}{2}$
Unrelated	0	1

Table 1: Genomic distances given by coefficient of kinships.

Similarity/Distances of Covariates In Section 3.3, we explored how genomes influence neuroanatomy by measuring brain volume, radial diffusivity (RD), axial diffusivity (AD), and fractional anisotropy (FA). These features were estimated per region of interest (ROI) defined by a parcellation, using the processed images obtained during network construction. RD is a measure of diffusion perpendicular to the axonal fibers, and it reflects the degree of damage or demyelination in the white matter. AD, on the other hand, is a measure of diffusion parallel to the axonal fibers and reflects the integrity and density of the axonal fibers. FA is a measure of the degree of anisotropy or directionality of water diffusion and reflects the coherence and organization of the white matter. We computed the average of these features across all regions, and we normalized the data by dividing each feature by its maximum value.

In order to perform the hypothesis test using DCORR in Section 3.3, we first compute the distance matrix for neuroanatomy. Let $w_i, w_j \in \mathbb{R}^d$ be measurements of observed covariates from subject i and j with dimensionality d. Then the distance function is simply given by the Euclidean norm:

$$d(w_i, w_j) = \|w_i - w_j\|_2$$

On the other hand, CDCORR requires a similarity matrix of the conditioning covariates, which includes the four neuroanatomy features and age. Age is normalized by dividing by the maximum age of the dataset. As before, let $w_i, w_j \in \mathbb{R}^d$ be measurements of observed covariates from subject i and j

with dimensionality d. The kernel function is

$$K(w_i, w_j) = (2\pi)^{\frac{1}{d}} |\mathbf{H}|^{\frac{1}{2}} \exp\left(\frac{1}{2}(w_i - w_j)^\top \mathbf{H}^{-2}(w_i - w_j)\right)$$

where $\mathbf{H}$ is a $d \times d$ diagonal matrix containing on the main diagonal the kernel bandwidth $(h_1, \ldots, h_d)$. While many methods exist for estimating the bandwidth, we chose a plug-in rule given by [73]:

$$h_i = n^{-1/(d+4)} \hat{\sigma}_i$$

where n is the number of observations, d is the dimensionality of each observation, and $\hat{\sigma}_i$ is the estimated standard deviation for dimension i.

5.5 p - values and Multiple Hypothesis Correction All p-values from DCORR and CDCORR tests are estimated using $N = 25{,}000$ permutations. Across all figures and tables associated with this work, we are concerned with obtaining a proper estimate of the rate at which we detect effects (*discoveries*). Therefore, we control the false discovery rate (FDR) with the Benjamini-Hochberg Correction [74].

6 Data and Code Availability Statement Analyses relied on graspologic [75], hyppo [76], NumPy [77], SciPy [78], Pandas [79], and NetworkX [80]. Plotting was performed using matplotlib [81] and Seaborn [82]. Processed connectomes can be found at s3://open-neurodata/hcp1200/. The data that support the findings of this study are available on request from the corresponding author JC. The data regarding subject-level data (e.g. age, family identifiers, sibling relationship identifiers) are restricted and not publicly available because the data might allow identification of individuals. All code used for this paper can be found at https://github.com/neurodata/connectomic-heritability and viewed as a JupyterBook [83] at http://docs.neurodata.io/connectomic-heritability.

7 Acknowledgments The authors are grateful for the support from the National Science Foundation (NSF) administered through NSF CAREER Award (Grant no. 1942963), the National Institute of Health (NIH) BRAIN Initiative (1RF1MH123233-01), and the National Institute of Health (NIH) through National Institute of Mental Health (NIMH) Research Project RF1MH128696-01. Data were provided [in part] by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

References

[1] Cristian Tomasetti and Bert Vogelstein. Variation in cancer risk among tissues can be explained by the number of stem cell divisions. Science, 347(6217):78–81, 2015.

[2] Naomi R Wray, Jian Yang, Michael E Goddard, and Peter M Visscher. The genetic interpretation of area under the roc curve in genomic profiling. PLoS genetics, 6(2):e1000864, 2010.

[3] Martin H Steinberg and Paola Sebastiani. Genetic modifiers of sickle cell disease. American journal of hematology, 87(8):795–803, 2012.

[4] Hongbo Xu, Qian Wang, Lei Yu, Aiju Liu, Jiayan Sheng, Tao Yang, and Hui Zhang. Crispr-cas9-based genome editing in hematological diseases: opportunities and challenges. Molecular Therapy-Methods & Clinical Development, 20:13–26, 2021.

[5] Jean Bennett, Jennifer Wellman, Kathleen A Marshall, Sarah McCague, Manzar Ashtari, Julie DiStefano-Pappas, Okan U Elci, Daniel C Chung, Junwei Sun, J Fraser Wright, et al. Gene therapy for hereditary blindness: Clinical trial of voretigene neparvovec. New England Journal of Medicine, 377(17):1644–1652, 2018.

[6] Punam Malik. Gene therapy for sickle cell disease: progress and challenges. Hematology–American Society of Hematology Education Program, 2020(1):9–16, 2020.

[7] Gergely L Lukacs. Gene therapy for cystic fibrosis: Progress and challenges. Molecular Therapy, 27(2):214–230, 2019.

[8] D.S. Falconer. Introduction to Quantitative Genetics. Always Learning. Pearson Education, 1996. ISBN 9788131727409.

[9] Ronald Aylmer Fisher. Statistical methods for research workers. In Breakthroughs in statistics, pages 66–70. Springer, 1992.

[10] Irving I Gottesman. Heritability of personality: a demonstration. Psychological Monographs: General and Applied, 77(9):1, 1963.

[11] Birgitta Floderus-Myrhed, Nancy Pedersen, and Ingrid Rasmuson. Assessment of heritability for personality, based on a short-form of the eysenck personality inventory: A study of 12,898 twin pairs. Behavior genetics, 10(2):153–162, 1980.

[12] Tena Vukasović and Denis Bratko. Heritability of personality: A meta-analysis of behavior genetic studies. Psychological bulletin, 141(4):769, 2015.

[13] Mark W Russell, Ian Law, Phyliss Sholinsky, and Richard R Fabsitz. Heritability of ecg measurements in adult male twins. Journal of electrocardiology, 30:64–68, 1998.

[14] JC Christian, PL Yu, CW Slemenda, and CC Johnston Jr. Heritability of bone mass: a longitudinal study in aging male twins. American journal of human genetics, 44(3):429, 1989.

[15] Wendy S Post, Martin G Larson, Richard H Myers, Maurizio Galderisi, and Daniel Levy. Heritability of left ventricular mass: the framingham heart study. Hypertension, 30(5):1025–1028, 1997.

[16] Bryan S Michalowicz, DP Aepplp, Ramesh K Kuba, JE Bereuter, JP Conry, NL Segal, TJ Bouchard Jr, and BL Pihlstrom. A twin study of genetic variation in proportional radiographic alveolar bone height. Journal of dental research, 70(11):1431–1435, 1991.

[17] Lindon J Eaves and John L Jinks. Insignificance of evidence for differences in heritability of iq between races and social classes. 1972.

[18] Philip A Vernon. The heritability of measures of speed of information-processing. Personality and Individual Differences, 10(5):573–576, 1989.

[19] Tyrone D Cannon, Jaakko Kaprio, Jouko Lönnqvist, Matti Huttunen, and Markku Koskenvuo. The genetic epidemiology of schizophrenia in a finnish twin cohort: a population-based modeling study. Archives of general psychiatry, 55(1):67–74, 1998.

[20] Isabelle F Foote, Benjamin M Jacobs, Georgina Mathlin, Cameron J Watson, Phazha LK Bothongo, Sheena Waters, Ruth Dobson, Alastair J Noyce, Kamaldeep S Bhui, Ania Korszun, et al. The genetic architecture of alzheimers disease risk: A genomic structural equation modelling study. medRxiv, 2021.

[21] Alexandra J Mayhew and David Meyre. Assessing the heritability of complex traits in humans: methodological challenges and opportunities. Current genomics, 18(4):332–340, 2017.

[22] MCCL Neale and Lon R Cardon. Methodology for genetic studies of twins and families, volume 67. Springer Science & Business Media, 2013.

[23] Martijn P. van den Heuvel, René C. W. Mandl, Cornelis J. Stam, René S. Kahn, and Hilleke E. Hulshoff Pol. Aberrant frontal and temporal complex network structure in schizophrenia: A graph theoretical analysis. Journal of Neuroscience, 30(47):15915–15926, Nov 2010. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.2874-10.2010.

[24] Qifeng Wang, Tung-Ping Su, Yuan Zhou, Kun-Hsien Chou, I-Yun Chen, Tianzi Jiang, and Ching-Po Lin. Anatomical insights into disrupted small-world networks in schizophrenia. Neuroimage, 59(2): 1085–1093, 2012.

[25] Sifis Micheloyannis. Graph-based network analysis in schizophrenia. World Journal of Psychiatry, 2(1):1–12, Feb 2012. ISSN 2220-3206. doi: 10.5498/wjp.v2.i1.1.

[26] Marc M. Bohlken, René C. W. Mandl, Rachel M. Brouwer, Martijn P. van den Heuvel, Anna M. Hedman, René S. Kahn, and Hilleke E. Hulshoff Pol. Heritability of structural brain network topology: A dti study of 156 twins. Human Brain Mapping, 35(10):5295–5305, 2014. ISSN 1097-0193. doi: 10.1002/hbm.22550.

[27] Benjamin Sinclair, Narelle K Hansell, Gabriëlla A M Blokland, Nicholas G Martin, Paul M Thompson, Michael Breakspear, Greig I de Zubicaray, Margaret J Wright, and Katie L McMahon. Heritability of the network architecture of intrinsic brain functional connectivity. Neuroimage, 121: 243–252, November 2015.

[28] Ed Bullmore and Olaf Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. Nature reviews neuroscience, 10(3):186–198, 2009.

[29] Sebastiano Galantucci, Federica Agosta, Elka Stefanova, Silvia Basaia, Martijn P Van Den Heuvel, Tanja Stojković, Elisa Canu, Iva Stanković, Vladana Spica, Massimiliano Copetti, et al. Structural brain connectome and cognitive impairment in parkinson disease. Radiology, 283(2):515–525, 2017.

[30] Hannelore Aerts, Wim Fias, Karen Caeyenberghs, and Daniele Marinazzo. Brain networks under attack: robustness properties and the impact of lesions. Brain, 139(12):3063–3083, 2016.

[31] Marc M Bohlken, René CW Mandl, Rachel M Brouwer, Martijn P van den Heuvel, Anna M Hedman, René S Kahn, and Hilleke E Hulshoff Pol. Heritability of structural brain network topology: a dti study of 156 twins. Human brain mapping, 35(10):5295–5305, 2014.

[32] Andrew E Reineberg, Alexander S Hatoum, John K Hewitt, Marie T Banich, and Naomi P Friedman. Genetic and environmental influence on the human functional connectome. Cerebral Cortex, 30(4):2099–2113, 2020.

[33] Jaewon Chung, Eric Bridgeford, Jesús Arroyo, Benjamin D Pedigo, Ali Saad-Eldin, Vivek Gopalakrishnan, Liang Xiang, Carey E Priebe, and Joshua T Vogelstein. Statistical connectomics. Annual Review of Statistics and Its Application, 8:463–492, 2021.

[34] Eric W Bridgeford, Jaewon Chung, Brian Gilbert, Sambit Panda, Adam Li, Cencheng Shen, Alexandra Badea, Brian Caffo, and Joshua T Vogelstein. Learning sources of variability from high-dimensional observational studies. arXiv preprint arXiv:2307.13868, 2023.

[35] Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. Biometrika, 70(1):41–55, 1983.

[36] Gábor J Székely, Maria L Rizzo, and Nail K Bakirov. Measuring and testing dependence by correlation of distances. The annals of statistics, 35(6):2769–2794, 2007.

[37] Gábor J Székely and Maria L Rizzo. Partial distance correlation with methods for dissimilarities. The Annals of Statistics, 42(6):2382–2412, 2014.

[38] Xueqin Wang, Wenliang Pan, Wenhao Hu, Yuan Tian, and Heping Zhang. Conditional distance correlation. Journal of the American Statistical Association, 110(512):1726–1734, 2015.

[39] Stephen J. Young and Edward R. Scheinerman. Random dot product graph models for social networks. page 138–149, 2007.

[40] Daniel L Sussman, Minh Tang, Donniell E Fishkind, and Carey E Priebe. A consistent adjacency spectral embedding for stochastic blockmodel graphs. Journal of the American Statistical Association, 107(499):1119–1128, 2012.

[41] Minh Tang, Avanti Athreya, Daniel L. Sussman, Vince Lyzinski, Youngser Park, and Carey E. Priebe. A semiparametric two-sample hypothesis testing problem for random graphs. Journal of Computational and Graphical Statistics, 26(2):344–354, 2017. doi: $10.1080/10618600.2016.1193505$. URL https://doi.org/10.1080/10618600.2016.1193505.

[42] Minh Tang, Avanti Athreya, Daniel L Sussman, Vince Lyzinski, and Carey E Priebe. A nonparametric two-sample hypothesis testing problem for random graphs. 2017.

[43] David C Van Essen, Stephen M Smith, Deanna M Barch, Timothy EJ Behrens, Essa Yacoub, Kamil Ugurbil, Wu-Minn HCP Consortium, et al. The wu-minn human connectome project: an overview. Neuroimage, 80:62–79, 2013.

[44] David C Van Essen, Kamil Ugurbil, E Auerbach, D Barch, TEJ Behrens, R Bucholz, Acer Chang, Liyong Chen, Maurizio Corbetta, Sandra W Curtiss, et al. The human connectome project: a data acquisition perspective. Neuroimage, 62(4):2222–2231, 2012.

[45] Gabriëlla AM Blokland, Greig I de Zubicaray, Katie L McMahon, and Margaret J Wright. Genetic and environmental influences on neuroimaging phenotypes: a meta-analytical perspective on twin imaging studies. Twin Research and Human Genetics, 15(3):351–371, 2012.

[46] Tian Ge, Martin Reuter, Anderson M Winkler, Avram J Holmes, Phil H Lee, Lee S Tirrell, Joshua L Roffman, Randy L Buckner, Jordan W Smoller, and Mert R Sabuncu. Multidimensional heritability analysis of neuroanatomical shape. Nature communications, 7(1):1–10, 2016.

[47] Rachel M Brouwer, Anna M Hedman, Neeltje EM van Haren, Hugo G Schnack, Rachel GH Brans, Dirk JA Smit, Rene S Kahn, Dorret I Boomsma, and Hilleke E Hulshoff Pol. Heritability of brain volume change and its relation to intelligence. Neuroimage, 100:676–683, 2014.

[48] Paul M Thompson, Tyrone D Cannon, Katherine L Narr, Theo Van Erp, Veli-Pekka Poutanen, Matti Huttunen, Jouko Lönnqvist, Carl-Gustaf Standertskjöld-Nordenstam, Jaakko Kaprio, Mohammad Khaledy, et al. Genetic influences on brain structure. Nature neuroscience, 4(12):1253–1258, 2001.

[49] Phil H Lee, Justin T Baker, Avram J Holmes, Neda Jahanshad, Tian Ge, Jae-Yoon Jung, Yanela Cruz, Dara S Manoach, Derrek P Hibar, Joshua Faskowitz, et al. Partitioning heritability analysis reveals a shared genetic basis of brain anatomy and schizophrenia. Molecular psychiatry, 21(12): 1680–1689, 2016.

[50] Anouk den Braber, Marc M Bohlken, Rachel M Brouwer, Dennis van't Ent, Ryota Kanai, René S Kahn, Eco JC de Geus, Hilleke E Hulshoff Pol, and Dorret I Boomsma. Heritability of subcortical brain measures: a perspective for future genome-wide association studies. NeuroImage, 83:98–102, 2013.

[51] Martijn P van den Heuvel and Olaf Sporns. A cross-disorder connectome landscape of brain dysconnectivity. Nature reviews neuroscience, 20(7):435–446, 2019.

[52] Alex Fornito, Andrew Zalesky, and Michael Breakspear. The connectomics of brain disorders. Nature Reviews Neuroscience, 16(3):159–172, 2015.

[53] Gregory Kiar, Eric Bridgeford, Will Gray Roncal, Vikram Chandrashekhar, Disa Mhembere, Sephira Ryman, Xi-Nian Zuo, Daniel S Marguiles, R Cameron Craddock, Carey E Priebe, Rex Jung, Vince Calhoun, Brian Caffo, Randal Burns, Michael P Milham, and Joshua Vogelstein. A high-throughput pipeline identifies robust connectomes but troublesome variability. bioRxiv, 2018. doi: $10.1101/188706$. URL https://www.biorxiv.org/content/early/2018/04/24/188706.

[54] Stephen M Smith et al. Advances in functional and structural MR image analysis and implementation as FSL. NeuroImage, 23 Suppl 1:S208–19, jan 2004. ISSN 1053-8119. URL

http://www.ncbi.nlm.nih.gov/pubmed/15501092.

[55] Mark W Woolrich et al. Bayesian analysis of neuroimaging data in FSL. NeuroImage, 45(1 Suppl): S173–86, mar 2009. ISSN 1095-9572. URL http://www.sciencedirect.com/science/article/pii/S1053811908012044.

[56] Mark Jenkinson et al. FSL. NeuroImage, 62(2):782–90, aug 2012. ISSN 1095-9572. URL http://www.ncbi.nlm.nih.gov/pubmed/21979382.

[57] John Mazziotta et al. A four-dimensional probabilistic atlas of the human brain. Journal of the American Medical Informatics Association, 8(5):401–430, 2001.

[58] Eleftherios Garyfallidis, Matthew Brett, Bagrat Amirbekian, Ariel Rokem, Stefan Van Der Walt, Maxime Descoteaux, and Ian Nimmo-Smith. Dipy, a library for the analysis of diffusion mri data. Frontiers in neuroinformatics, 8:8, 2014.

[59] Alexander Schaefer, Ru Kong, Evan M Gordon, Timothy O Laumann, Xi-Nian Zuo, Avram J Holmes, Simon B Eickhoff, and BT Thomas Yeo. Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri. Cerebral cortex, 28(9):3095–3114, 2018.

[60] Nathalie Tzourio-Mazoyer, Brigitte Landeau, Dimitri Papathanassiou, Fabrice Crivello, Olivier Etard, Nicolas Delcroix, Bernard Mazoyer, and Marc Joliot. Automated anatomical labeling of activations in spm using a macroscopic anatomical parcellation of the mni mri single-subject brain. Neuroimage, 15(1):273–289, 2002.

[61] 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, et al. A multimodal parcellation of human cerebral cortex. Nature, 536(7615):171–178, 2016.

[62] Rahul S Desikan, Florent Ségonne, Bruce Fischl, Brian T Quinn, Bradford C Dickerson, Deborah Blacker, Randy L Buckner, Anders M Dale, R Paul Maguire, Bradley T Hyman, et al. An automated labeling system for subdividing the human cerebral cortex on mri scans into gyral based regions of interest. Neuroimage, 31(3):968–980, 2006.

[63] Ross M. Lawrence, Eric W. Bridgeford, Patrick E. Myers, Ganesh C. Arvapalli, Sandhya C. Ramachandran, Derek A. Pisner, Paige F. Frank, Allison D. Lemmer, Aki Nikolaidis, and Joshua T. Vogelstein. Standardizing human brain parcellations. Sci. Data, 8(78):1–9, March 2021. ISSN 2052-4463. doi: 10.1038/s41597-021-00849-3.

[64] Anna Goldenberg, Alice X. Zheng, Stephen E. Fienberg, and Edoardo M. Airoldi. A survey of statistical network models. Found. Trends Mach. Learn., 2(2):129–233, February 2010. ISSN 1935-8237. doi: 10.1561/2200000005. URL http://dx.doi.org/10.1561/2200000005.

[65] Eric D Kolaczyk and Gábor Csárdi. Statistical analysis of network data with R, volume 65. Springer, 2014.

[66] Avanti Athreya, Donniell E. Fishkind, Minh Tang, Carey E. Priebe, Youngser Park, Joshua T. Vogelstein, Keith Levin, Vince Lyzinski, Yichen Qin, and Daniel L Sussman. Statistical inference on random dot product graphs: a survey. Journal of Machine Learning Research, 18(226):1–92, 2018. URL http://jmlr.org/papers/v18/17-448.html.

[67] Daniel L. Sussman, Minh Tang, Donniell E. Fishkind, and Carey E. Priebe. A consistent adjacency spectral embedding for stochastic blockmodel graphs. Journal of the American Statistical Association, 107(499):1119–1128, 2012. doi: 10.1080/01621459.2012.699795. URL https://doi.org/10.1080/01621459.2012.699795.

[68] Daniel L. Sussman, Minh Tang, and Carey E. Priebe. Consistent latent position estimation and vertex classification for random dot product graphs. IEEE Transactions on Pattern Analysis and Machine Intelligence, 36:48–57, 2014. doi: 10.1109/TPAMI.2013.135.

[69] Vince Lyzinski, Daniel L. Sussman, Minh Tang, Avanti Athreya, and Carey E. Priebe. Perfect clustering for stochastic blockmodel graphs via adjacency spectral embedding. Electronic Journal of Statistics, 8(2):2905–2922, 2014. doi: 10.1214/14-EJS978. URL https://doi.org/10.1214/14-EJS978.

[70] Mu Zhu and Ali Ghodsi. Automatic dimensionality selection from the scree plot via the use of profile likelihood. Computational Statistics & Data Analysis, 51(2):918–930, 2006.

[71] Runze Tang, Michael Ketcha, Alexandra Badea, Evan D Calabrese, Daniel S Margulies, Joshua T Vogelstein, Carey E Priebe, and Daniel L Sussman. Connectome smoothing via low-rank approximations. IEEE transactions on medical imaging, 38(6):1446–1456, 2018.

[72] Doug Speed, Gibran Hemani, Michael R Johnson, and David J Balding. Improved heritability estimation from genome-wide snps. The American Journal of Human Genetics, 91(6):1011–1021, 2012.

[73] David W Scott. Multivariate density estimation: theory, practice, and visualization. John Wiley & Sons, 2015.

[74] Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. Journal of the Royal statistical society: series B (Methodological), 57 (1):289–300, 1995.

[75] Jaewon Chung, Benjamin D Pedigo, Eric W Bridgeford, Bijan K Varjavand, Hayden S Helm, and Joshua T Vogelstein. GraSPy: Graph Statistics in Python. J. Mach. Learn. Res., 20(158):7, 2019.

[76] Sambit Panda, Satish Palaniappan, Junhao Xiong, Eric W. Bridgeford, Ronak Mehta, Cencheng Shen, and Joshua T. Vogelstein. Hyppo: A Comprehensive Multivariate Hypothesis Testing Python Package. arXiv:1907.02088 [cs, stat], August 2020.

[77] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. Nature, 585(7825):357–362, September 2020. ISSN 1476-4687. doi: 10.1038/s41586-020-2649-2. URL https://www.nature.com/articles/s41586-020-2649-2. Number: 7825 Publisher: Nature Publishing Group.

[78] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C. J. Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, and Paul van Mulbregt. SciPy 1.0: fundamental algorithms for scientific computing in Python. Nature Methods, 17(3):261–272, March 2020. ISSN 1548-7105. doi: 10.1038/s41592-019-0686-2. URL https://www.nature.com/articles/s41592-019-0686-2. Number: 3 Publisher: Nature Publishing Group.

[79] Wes McKinney. pandas: a Foundational Python Library for Data Analysis and Statistics. Python for high performance and scientific computing, 14(9):1–9, 2011. Publisher: Seattle.

[80] Aric Hagberg, Pieter Swart, and Daniel Schult. Exploring network structure, dynamics, and function using networkx. Technical Report LA-UR-08-05495; LA-UR-08-5495, Los Alamos National Lab. (LANL), Los Alamos, NM (United States), January 2008. URL https://www.osti.gov/biblio/960616.

[81] John D. Hunter. Matplotlib: A 2D Graphics Environment. Computing in Science & Engineering, 9(3):90–95, May 2007. ISSN 1558-366X. doi: 10.1109/MCSE.2007.55. Conference Name: Computing in Science & Engineering.

[82] Michael L. Waskom. seaborn: statistical data visualization. Journal of Open Source Software, 6 (60):3021, 2021. doi: 10.21105/joss.03021. URL https://doi.org/10.21105/joss.03021. Publisher: The Open Journal.

[83] Executable Books Community. Jupyter book, February 2020. URL https://doi.org/10.5281/zenodo

.4539666.