
Angle Basis: a Generative Model and Decomposition for Functional Connectivity

Anton Orlichenko*

Department of Biomedical Engineering
Tulane University
New Orleans, LA 70118
aorlichenko@tulane.edu

Gang Qu

Department of Biomedical Engineering
Tulane University
New Orleans, LA 70118
gqu1@tulane.edu

Ziyu Zhou

Department of Computer Science
Tulane University
New Orleans, LA 70118
zzhou11@tulane.edu

Zhengming Ding

Department of Computer Science
Tulane University
New Orleans, LA 70118
zding1@tulane.edu

Yu-Ping Wang

Department of Biomedical Engineering
Tulane University
New Orleans, LA 70118
wyp@tulane.edu

Abstract

Functional connectivity (FC) is one of the most common inputs to fMRI-based predictive models, due to a combination of its simplicity and robustness. However, there may be a lack of theoretical models for the generation of FC. In this work, we present a straightforward decomposition of FC into a set of basis states of sine waves with an additional jitter component. We show that the decomposition matches the predictive ability of FC after including 5-10 bases. We also find that both the decomposition and its residual have approximately equal predictive value, and when combined into an ensemble, exceed the AUC of FC-based prediction by up to 5%. Additionally, we find the residual can be used for subject fingerprinting, with 97.3% same-subject, different-scan identifiability, compared to 62.5% for FC. Unlike PCA or Factor Analysis methods, our method does not require knowledge of a population to perform its decomposition; a single subject is enough. Our decomposition of FC into two equally-predictive components may lead to a novel appreciation of group differences in patient populations. Additionally, we generate synthetic patient FC based on user-specified characteristics such as age, sex, and disease diagnosis. By creating synthetic datasets or augmentations we may reduce the high financial burden associated with fMRI data acquisition.

1 Introduction and Related Work

Functional magnetic resonance (fMRI) uses the blood oxygen level-dependent (BOLD) signal to map blood flow in the brain [2], which is correlated to neural activity through the hemodynamic response function [39]. Early work in fMRI utilized statistical techniques on voxel intensities to

*Corresponding author.

measure brain region activation and group differences [32]. More recently, functional connectivity (FC), the temporal Pearson correlation of BOLD signal across different regions of the brain, has been used to characterize brain state [16]. Regions are delineated using either a template [37] or ICA [8].

FC has recently been used for the prediction of subject age [25], sex [50], race [33], or schizophrenia (SZ) diagnosis [24], at levels that in some cases meet or exceed prediction using genetic information [5]. It has also been used to flag possible pre-clinical Alzheimer’s disease by predicting a subject’s brain age based on FC and comparing with actual age [30]. Interestingly, it has even been used for mechanistic studies of aggression in response to olfactory stimulus [31], the traditional province of statistical fMRI techniques. In the hospital setting, it may be used in evaluating the severity and monitoring the treatment of concussions [18]. Dynamic FC, using either window-based or non-window-based [4] methods, has also been explored in the literature, with some studies showing benefit over static FC [40] and others disagreeing [23].

Many studies have explored alternative measures of connectivity, such precision matrix-based (partial) correlation [20], distance correlation [13][28], effective connectivity [43], or phase lock value (PLV) [14]. In general, these methods have been found to be largely similar [1] in terms of usefulness. For example, partial correlation may have better test-retest reliability, but lose the meaning of region-region synchronization inherent in FC [29]. Geerligts et al. [13] found that Pearson correlation has slightly worse results compared to distance correlation. In fact, FC and distance correlation-based connectivity maps are largely indistinguishable [28]. BrainNetCNN [19], by contrast, works directly on time series data, but the resulting convolutional filters are not as interpretable as FC maps. Since FC is an established method, many recent papers focus on novel approaches. It should be noted, however, that even since the days of ANOVA-based statistical methods, fMRI has suffered from a reproducibility crisis [3][10]. Therefore a focus on simplicity and reproducibility [48] may be preferred.

In terms of generative models which aim to simulate brain dynamics, the main focus has been on structural equation modeling (SEM) or dynamic causal modeling (DCM), used to generate effective connectivity [26]. Unfortunately, DCM has a high computational complexity that makes it unsuitable for working with a large number of regions [17][41]. Additionally, even the highest resolution fMRI only identifies the averaged activity of many individual neurons, further blurred by the hemodynamic response function [39]. Graphical models use FC, or some other connectivity method, as the basis for determining graph edge weights and then use thresholding to binarize these edges [21]. In effect, FC is the foundation on which most such methods are built.

Zhao et al. [51] and others [27] have used GANs to help discriminate between patients and healthy controls. Tan et al. [47] have used a manifold-regularized Wasserstein distance GAN to achieve modest prediction gains. However, these GAN-based methods fail to model the underlying time-varying BOLD signal that underpins FC. Therefore they cannot be used for decomposition, dimensionality reduction, or easy interpretation as can our work.

In this work, we provide a theoretical basis for decomposition of FC based on regions of in-phase or out-of-phase sine waves, representing the bandpass-filtered BOLD signal at a region of interest (ROI). In conceptual terms, it has some similarity to the PLV determined using the Hilbert transform [14]. The PLV, however, is not widely used in predictive models [49], and our own validation experiments have shown it performs much worse than other measures of connectivity when used for prediction tasks (see Supplemental Materials). By combining a sine-wave decomposition with a jitter value at each ROI, we are able to reconstruct the majority of FC, creating both a useful decomposition and a useful residual. Typical MRI studies require \$500-\$1000 per subject for data acquisition [45]. Using our model for the generation of synthetic FC based on specific phenotype settings, we may alleviate some of the high financial burden of fMRI data acquisition, allowing greater accessibility to this modality.

2 Methods

We describe angle basis as an approximation of the Fourier series representation of the time domain BOLD fMRI signal $x(t)$, from which FC is derived. We motivate the underlying model by a

connection to the well-known but less widely used PLV. Note that the FC is defined as the Pearson correlation of the time domain signals between regions c and d :

$$\rho_{cd} = \frac{\sigma_{cd}^2}{\sqrt{\sigma_c^2 \sigma_d^2}} \quad (1)$$

2.1 Phase Lock Value (PLV)

Any time domain signal can be represented in a basis of complex sinusoids via the Fourier transform:

$$X(\omega) = \int_{-\infty}^{\infty} x(t) e^{-i\omega t} dt \quad (2)$$

The Hilbert transform, acting on the time domain signal, imparts a $\pi/2$ phase shift to every complex sinusoid depending on the sign of the sinusoid frequency:

$$\mathbf{H}[x](t) = \frac{1}{\pi} \text{p.v.} \int_{-\infty}^{\infty} \frac{x(\tau)}{t - \tau} d\tau \quad (3)$$

Here p.v. represents the Cauchy principal value of the improper integral. Using the Hilbert transform, one may change the original signal into an analytic signal $x_a(t)$ containing only positive frequencies.

$$x_a(t) = x(t) + i\mathbf{H}[x(t)] \quad (4)$$

The analytic signal may be factorized into a low-frequency amplitude envelope $a(t)$ and a phase term $\theta(t)$, so long as Bedrosian's theorem [6] holds, i.e. as long as the support of $a(t)$ and $\theta(t)$ are different. This may be the case for fMRI signals, which are typically bandpass filtered.

$$x_a(t) = a(t) e^{i\theta(t)} \quad (5)$$

The connectivity between two regions c, d may be quantified as the PLV, a measure of the dispersion of phase differences between the two regions [14].

$$\begin{aligned} \theta_{cd}(t) &= \theta_c(t) - \theta_d(t) \\ \text{PLV}_{cd} &= \frac{1}{T} \left| \sum_{t=1}^T e^{i\theta_{cd}(t)} \right| \end{aligned} \quad (6)$$

Although PLV has been widely used in neuroscience [22], it is dependent on the passband and has relatively poor predictive ability (see Supplemental Materials). Additionally, one may wish to work directly with the FC. Given this motivation, we present a simplified but more robust model.

2.2 Angle Basis Decomposition

First, we represent the signal as a Fourier series, assuming, as in the case of the PLV, that the amplitude envelope of the passband-filtered signal is constant with respect to the phase signal.

$$x(t) = \sum_{n=1}^N A_n \sin(\omega_n t + \theta_n) \quad (7)$$

In the simplest case, we approximate the BOLD signal at each region c as a single sine wave with arbitrary constant amplitude and phase.

$$\hat{x}_c(t) = A_c \sin(\omega t + \theta_c) \quad (8)$$

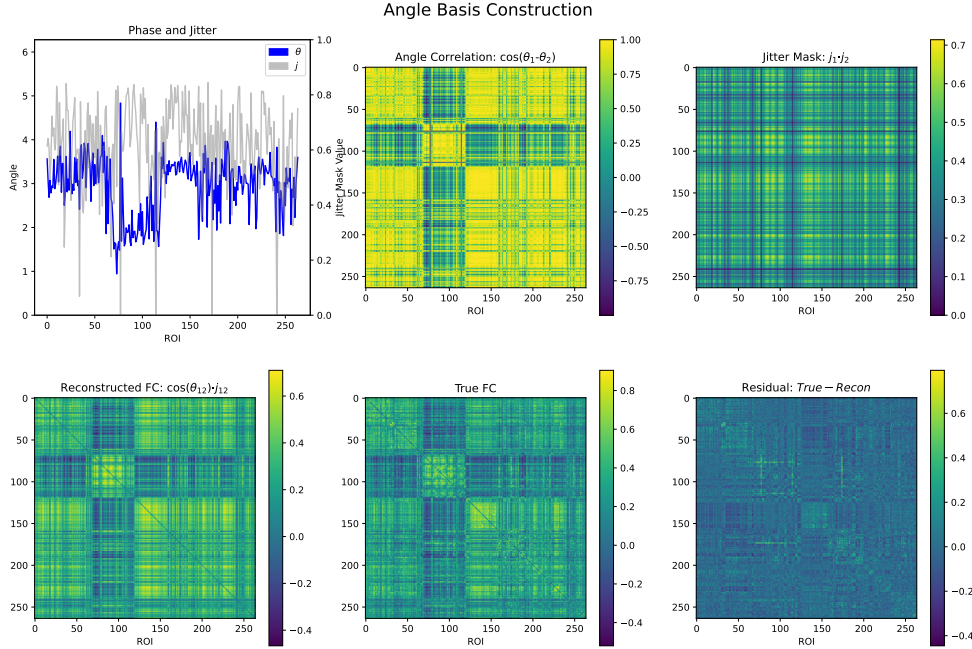


Figure 1: Example of the angle basis reconstruction of one representative subject from the PNC dataset. The prominent block corresponds to the default mode network (DMN). The scan is probably of a female, due to the high within-DMN connectivity [11]. Here only one basis is used.

The correlation between ROIs c and d is then simply the cosine of the phase difference.

$$\hat{\rho}_{cd} = \cos(\theta_c - \theta_d) \quad (9)$$

We reduce the number of frequencies required to decompose the signal by introducing a jitter component, which models the signal at a region "jittering" such that it becomes uncorrelated with all other regions.

$$\hat{\rho}_{cd}^{(n)} = j_c^{(n)} \cdot j_d^{(n)} \cdot \cos(\theta_c^{(n)} - \theta_d^{(n)}) \quad (10)$$

Here the superscript (n) refers to the correlation in the n^{th} basis. Each scan can be modeled with up to N bases, where the bases represent either different temporal regions or orthogonal sinusoids of different frequencies. We then calculate the overall correlation between regions c, d as follows:

$$\tilde{\rho}_{cd} = \frac{1}{N} \sum_{n=1}^N \hat{\rho}_{cd}^{(n)} \quad (11)$$

We understand this formula is only valid when amplitudes are the same across all basis sinusoids; however we find this model to be useful regardless. Additionally, differences in window lengths or amplitudes can be incorporated into the $j_c^{(n)} \cdot j_d^{(n)}$ term in a uniform way.

The decomposition is estimated via gradient descent by minimizing the RMSE between the reconstructed and true FC, i.e. the residual r_{cd} .

$$\min_{\theta, j} \frac{1}{R^2} \sum_{cd} r_{cd}^2 \quad \text{s.t.} \quad 0 \leq \theta_c^{(n)} < 2\pi, \quad 0 \leq j_c^{(n)} \leq 1 \quad (12)$$

Here R is the number of regions in the remplate, $\rho_{cd}, \tilde{\rho}_{cd}$ are the true and estimated FC values, respectively, and r_{cd} is the residual. Values of jitter that become less than 0 or greater than 1 are clipped to those values while phase angles wrap around.

Figure 1 shows the decomposition of a representative’s subject’s FC into a single basis plus residual. The template we used contains 264 regions, meaning that each basis must contain 528 values for phase and jitter. In section 3.4, we show that 5-10 bases are enough to recreate the predictive capacity of FC. In contrast, there are 34,716 non-duplicate values in a subject’s FC.

Angle basis estimation is carried out using PyTorch [34] with CUDA support.² Our source code may be found at <https://github.com/aorliche/AngleBasis>. Training time is approximately one second per subject for a one basis decomposition, up to 10 seconds for higher numbers of basis sinusoids, using a single GPU.

3 Experiments

We validate the usefulness of the decomposition in five downstream tasks:

- Identification of the same subject from different scans
- Prediction of age, sex, race, and schizophrenia diagnosis across two different datasets
- Transfer of models trained on one dataset to predict sex or race to another dataset
- Prediction using reduced feature numbers (feature selection)
- Generation of synthetic FC

Prediction was carried out using vectorizations of the true FC matrix, angle basis reconstruction, angle basis residual, or ensemble of angle basis and residual. In addition, a Factor Analysis was used to create latents and residuals for the identifiability task and age prediction. Identifiability was carried out by cosine similarity, which was found to give better results than Euclidean distance for all inputs tested. A deep autoencoder (AE) with 2 hidden layers of 2000 and 2000 neurons with ReLU activation was also used to create a reconstruction/residual for the prediction tasks.

The models used for prediction were simple Ridge and Logistic Regression models, from the scikit-learn implementation[35],³ with a maximum of 1000 iterations and regularization parameters $\alpha = 1$ and $C = 1$, respectively. Models were evaluated using an 80/20 train test split over 20 bootstrap repetitions. When using Ridge Regression (age prediction), both training and test sets were centered using the mean age of the training set. For race prediction, only African Ancestry (AA) and European Ancestry (EA) ethnicities were used, because our datasets were enriched for these groups, with other ethnicities making up less than 10% of subjects. Although FC-based predictions are not always improved by using more expressive models [36], we include a comparison with an MLP and BrainNetCNN [19] in the Supplemental Materials.

3.1 Datasets

We tested our model on two different, large, widely-used datasets available to researchers via application. Both datasets were pre-processed with SPM12,⁴ including intra-subject co-registration, warping to MNI space, smoothing with a 5mm FWHM kernel, temporal bandpass filtering from 0.01 to 0.15 Hz, and time series extraction using the 264-region Power atlas [37]. Power atlas regions are discrete balls of 5mm radius from which average BOLD signal was extracted.

²<https://pytorch.org/>

³<https://scikit-learn.org/stable/>

⁴<https://www.fil.ion.ucl.ac.uk/spm/software/spm12/>

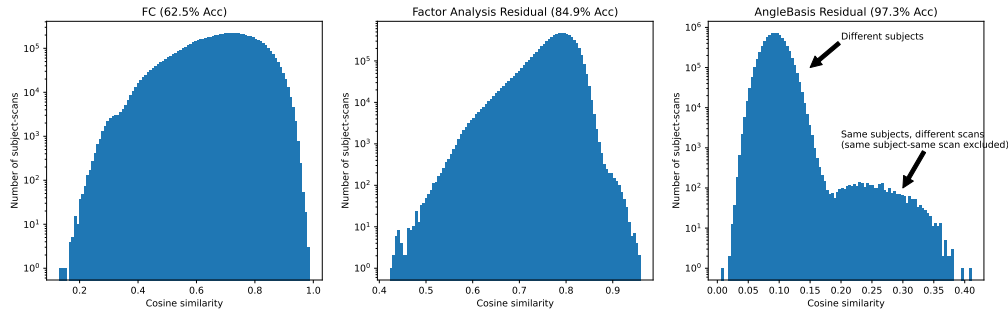


Figure 2: Identification of subjects from different scans in the PNC dataset. Histograms show the number of scan pairs at a particular cosine similarity value. Most subjects in the PNC dataset contain 3 scans with different in-scanner tasks. A total of 3849 scans were used.

3.1.1 Philadelphia Neurodevelopmental Cohort (PNC)

The PNC dataset is a large, widely used dataset consisting of children and young adults 8-22 years old, 1,529 of whom have fMRI scans [42] and over 9,000 of whom have genetic information [15]. Most subjects have three scans, one for a resting state, one for a working memory [38], and one for an emotion identification in-scanner task. A wide variety of phenotype information is available, including sex, race, cognitive battery, and questionnaire info [9]. Subjects were not specifically selected due to prior neurological problems. Of subjects having scans, 725 are male, 804 are female, 660 are AA ethnicity, 690 are EA ethnicity, and 179 are of other ethnicity. More detailed phenotype data and correlations are provided in the Supplemental Materials.

3.1.2 Bipolar and Schizophrenia Network for Intermediate Phenotypes (BSNIP)

The BSNIP dataset is a large multi-site fMRI dataset acquired to further understand schizophrenia, schizoaffective disorder, and bipolar disorder [46]. Our version of the dataset contains fMRI scans for 1,245 adult subjects acquired over 6 sites. These include 509 total patients (of which 199 have SZ diagnosis), 494 relatives of patients, and 242 normal controls. Subjects demographics are as follows: 528 are male, 717 are female, 387 are AA ethnicity, 778 are Caucasian (CA) ethnicity, and 81 are of other ethnicity. See Supplemental Materials for additional phenotype data.

3.2 Identification of Individuals

We find that taking the residual of the angle basis reconstruction using one basis greatly increases the identification rate of subjects in the PNC dataset, from 62.5% using FC to 97.3% using the angle basis residual (see Figure 2). In addition, we find that a Factor Analysis residual (not PCA) with $N = 10$ factors also increases identification rate to 84.9%. Previous work has found similar different-task same-subject identification as our FC result [12], as well as a slightly better result using a Factor Analysis-like method to produce a residual [7]. However, as seen in Figure 3, we find that the Factor Analysis method is not amenable to high prediction accuracy. Additionally, Factor Analysis, PCA, and the deep AE require knowledge of and have a dependence on the entire dataset.

3.3 Prediction Accuracy, Model Transfer, and Feature Selection

Figure 3 shows that prediction RMSE of age in the PNC dataset is comparable between FC and the first residual, and slightly better for the ensemble of residual and angle basis reconstruction. In addition, we see that removing the confounding effects of sex and race results in much stronger age prediction accuracy. In fact, increasing the dataset by a factor of four still doesn't match the accuracy of predicting within sex and race-matched groups. We also note that the Factor Analysis method does not have strong predictive fidelity for age prediction, and the deep AE ensemble performs poorly in predictive tasks. PCA was similar to the deep AE (not shown). We additionally tested the prediction fidelity of sex and race in the PNC dataset; and sex, race, and schizophrenia diagnosis in the BSNIP dataset. The results are shown in Table 1. That the residual is highly useful for prediction is supported by related work manipulating PCA-based identifiability to enhance predictive performance [44].

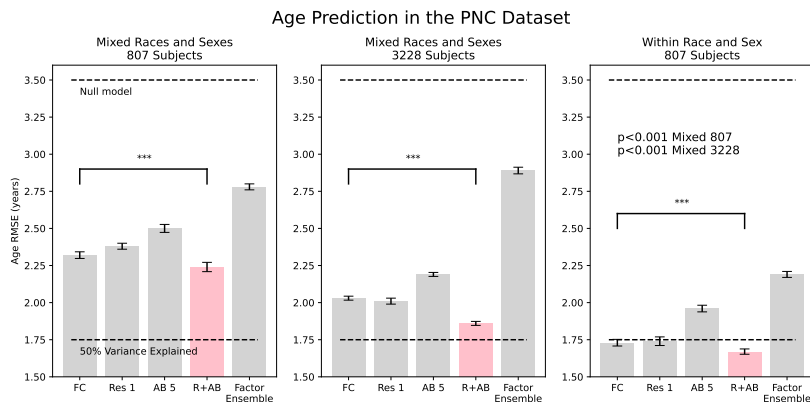


Figure 3: Age prediction in the PNC dataset. The ensemble of residual and angle basis is slightly but significantly better than the FC from which it was derived. Additionally, we find a strong effect for predicting within race and sex-matched compared to within mixed groups.

Table 1: Classification Accuracy (AUC)

Dataset	Predictive Task	FC	Deep AE Ens	AB+Res Ens	p-value
BSNIP	SZ/NC	0.785 ± 0.048	0.701 ± 0.062	0.804 ± 0.040	0.240
BSNIP	Sex	0.755 ± 0.023	0.645 ± 0.085	0.791 ± 0.022	0.002
BSNIP	Race	0.845 ± 0.022	0.739 ± 0.051	0.866 ± 0.023	0.108
PNC	Sex	0.886 ± 0.006	0.744 ± 0.065	0.923 ± 0.010	< 0.001
PNC	Race	0.946 ± 0.007	0.812 ± 0.040	0.973 ± 0.003	< 0.001
BSNIP→PNC	Sex	0.667 ± 0.017	0.601 ± 0.032	0.700 ± 0.013	< 0.001
BSNIP→PNC	Race	0.807 ± 0.018	0.710 ± 0.010	0.847 ± 0.012	< 0.001
PNC→BSNIP	Sex	0.629 ± 0.019	0.572 ± 0.019	0.667 ± 0.013	< 0.001
PNC→BSNIP	Race	0.800 ± 0.010	0.702 ± 0.022	0.832 ± 0.009	< 0.001

We next tested the transfer of sex and race prediction models from the PNC to the BSNIP dataset and vice versa. Models were trained wholly on one dataset and evaluated on the other. The results are shown in Table 1. It should be remembered that the PNC dataset is composed of ostensibly normal pre-teens, teens, and young adults, while the BSNIP dataset is composed of adults with schizophrenia-like diseases, their relatives, and normal controls. The high model transfer AUC for race prediction makes the case for biologically significant invariant features. Sex, by contrast, is less stable to model transfer, though prediction is still somewhat robust. It is well-known, for instance, that, on average, women demonstrate stronger within-DMN connectivity than men [11].

We used both Ridge and Lasso models to estimate the most relevant features for age prediction in the PNC dataset. Features were estimated on 50% of the dataset, with the remaining 50% used in an 80/20 split to train on and validate the selected features. See Supplemental Materials for results.

3.4 Effect of the Number of Bases

We find that prediction accuracy increases monotonically with increasing number of bases (see Figure 4), corresponding to decreasing reconstruction error. In contrast, prediction accuracy of the residual decreases monotonically with increasing number of bases. Using the first residual and the 20th angle basis reconstruction in an ensemble results in a higher accuracy than FC. Using FC twice in an ensemble is not significantly different compared to a single FC-based prediction.

3.5 Generation of Synthetic FC

We trained a set of 20 predictive models for age, race, sex, and schizophrenia diagnosis, and used them to generate synthetic FC via the angle basis framework using random seeds. The results for a

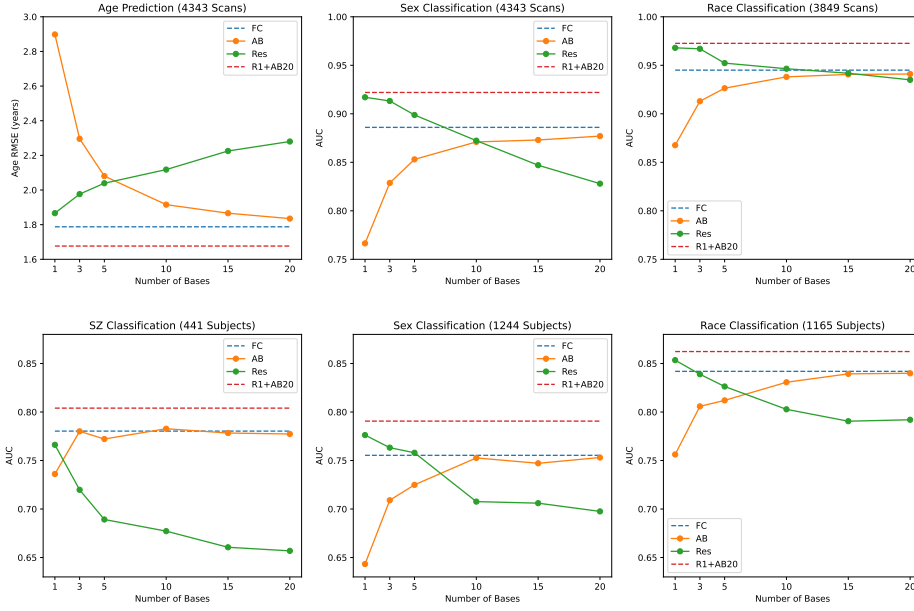


Figure 4: The prediction fidelity (RMSE and AUC) of the angle basis reconstruction and residual as a function of the number of bases. Top plot is for the PNC dataset, bottom is for the BSNIP dataset.

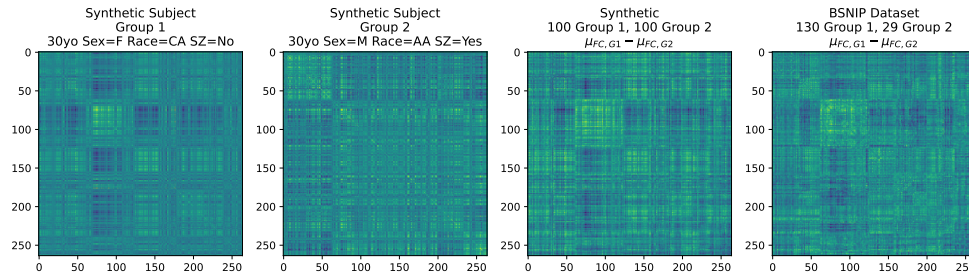


Figure 5: Generation of synthetic subjects that capture characteristics of patient FC. The synthetic female usually, but not always, has strong intra-DMN connectivity, while the synthetic man has a more flat connectivity profile. See the Supplemental Materials for additional examples and the distribution of synthetic scans for a given phenotype input.

synthetic 30 year old normal CA female and synthetic 30 year old schizophrenic AA male are shown in Figure 5. Additional examples and the procedure for generating synthetic FC are given in the Supplemental Materials.

4 Conclusion

In this work, we present a generative model for and decomposition of FC, which requires only a single subject’s connectivity matrix as input, and which generates a highly accurate reconstruction using a fraction of the original parameters. We find that both the reconstruction and residual have similar predictive ability, with an ensemble of the two surpassing FC in most cases. Additionally, we find that the residual is highly useful in subject identification or fingerprinting, achieving a nearly 100% subject identification rate on different in-scanner task scans. The estimation of the decomposition is fast, taking seconds for a single subject using a workstation with a mid-priced GPU. Our goal is to use this work as a baseline for the creation of synthetic FC datasets that incorporate information about population characteristics, with the aim to reduce the financial burden of fMRI acquisition.

5 Acknowledgements

The authors would like to acknowledge the NIH (grants R01 GM109068, R01 MH104680, R01 MH107354, P20 GM103472, R01 REB020407, R01 EB006841, R56 MH124925) and NSF (#1539067) for partial funding support.

fMRI and phenotype data came from the Neurodevelopmental Genomics: Trajectories of Complex Phenotypes database of genotypes and phenotypes repository, dbGaP Study Accession ID phs000607.v3.p2, as well as the Bipolar and Schizophrenia Network for Intermediate Phenotypes study (<http://b-snip.org/>).

References

- [1] H. Ahmadi, E. Fatemizadeh, and A. Motie-Nasrabadi. A comparative study of correlation methods in functional connectivity analysis using fMRI data of Alzheimer's patients. *J. Biomed. Phys. Eng.*, 13(2):125–134, Apr. 2023.
- [2] J. W. Belliveau, D. N. Kennedy, R. C. McKinstry, B. R. Buchbinder, R. M. Weisskoff, M. S. Cohen, J. M. Vevea, T. J. Brady, and B. R. Rosen. Functional mapping of the human visual cortex by magnetic resonance imaging. *Science*, 254 5032:716–9, 1991.
- [3] C. M. Bennett and M. B. Miller. How reliable are the results from functional magnetic resonance imaging? *Ann. N. Y. Acad. Sci.*, 1191(1):133–155, Mar. 2010.
- [4] T. Bolton and D. Van De Ville. Sparse coupled hidden Markov models shed light on resting-state fMRI cross-network interactions. In *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*. IEEE, Apr. 2017.
- [5] M. Bracher-Smith, K. Crawford, and V. Escott-Price. Machine learning for genetic prediction of psychiatric disorders: a systematic review. *Mol. Psychiatry*, 26(1):70–79, Jan. 2021.
- [6] J. L. Brown. A Hilbert transform product theorem. *Proc. IEEE Inst. Electr. Electron. Eng.*, 74(3):520–521, 1986.
- [7] B. Cai, G. Zhang, W. Hu, A. Zhang, P. Zille, Y. Zhang, J. M. Stephen, T. W. Wilson, V. D. Calhoun, and Y.-P. Wang. Refined measure of functional connectomes for improved identifiability and prediction. *Hum. Brain Mapp.*, 40(16):4843–4858, Nov. 2019.
- [8] V. D. Calhoun, J. Liu, and T. Adali. A review of group ICA for fMRI data and ICA for joint inference of imaging, genetic, and ERP data. *Neuroimage*, 45(1 Suppl):S163–72, Mar. 2009.
- [9] M. E. Calkins, K. R. Merikangas, T. M. Moore, M. Burstein, M. A. Behr, T. D. Satterthwaite, K. Ruparel, D. H. Wolf, D. R. Roalf, F. D. Mentch, H. Qiu, R. Chiavacci, J. J. Connolly, P. M. A. Sleiman, R. C. Gur, H. Hakonarson, and R. E. Gur. The Philadelphia neurodevelopmental cohort: constructing a deep phenotyping collaborative. *J. Child Psychol. Psychiatry*, 56(12):1356–1369, Dec. 2015.
- [10] A. Eklund, T. E. Nichols, and H. Knutsson. Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates. *Proc. Natl. Acad. Sci. U. S. A.*, 113(28):7900–7905, July 2016.
- [11] B. Ficek-Tani, C. Horien, S. Ju, W. Xu, N. Li, C. Lacadie, X. Shen, D. Scheinos, T. Constable, and C. Fredericks. Sex differences in default mode network connectivity in healthy aging adults. *Cereb. Cortex*, Dec. 2022.
- [12] E. S. Finn, X. Shen, D. Scheinost, M. D. Rosenberg, J. Huang, M. M. Chun, X. Papademetris, and R. T. Constable. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nat. Neurosci.*, 18(11):1664–1671, Nov. 2015.
- [13] L. Geerligs, Cam-Can, and R. N. Henson. Functional connectivity and structural covariance between regions of interest can be measured more accurately using multivariate distance correlation. *Neuroimage*, 135:16–31, July 2016.
- [14] E. Glerean, J. Salmi, J. M. Lahnakoski, I. P. Jääskeläinen, and M. Sams. Functional magnetic resonance imaging phase synchronization as a measure of dynamic functional connectivity. *Brain Connect.*, 2(2):91–101, June 2012.

- [15] J. T. Glessner, M. P. Reilly, C. E. Kim, N. Takahashi, A. Albano, C. Hou, J. P. Bradfield, H. Zhang, P. M. A. Sleiman, J. H. Flory, M. Imielinski, E. C. Frackelton, R. Chiavacci, K. A. Thomas, M. Garris, F. G. Otieno, M. Davidson, M. Weiser, A. Reichenberg, K. L. Davis, J. I. Friedman, T. P. Cappola, K. B. Margulies, D. J. Rader, S. F. A. Grant, J. D. Buxbaum, R. E. Gur, and H. Hakonarson. Strong synaptic transmission impact by copy number variations in schizophrenia. *Proc. Natl. Acad. Sci. U. S. A.*, 107(23):10584–10589, June 2010.
- [16] M. D. Greicius, B. Krasnow, A. L. Reiss, and V. Menon. Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences*, 100(1):253–258, 2003. doi: 10.1073/pnas.0135058100. URL <https://www.pnas.org/doi/abs/10.1073/pnas.0135058100>.
- [17] E. Hidalgo-Lopez, P. Zeidman, T. Harris, A. Razi, and B. Pletzer. Spectral dynamic causal modelling in healthy women reveals brain connectivity changes along the menstrual cycle. *Commun. Biol.*, 4(1):954, Aug. 2021.
- [18] M. Kaushal, L. Y. España, A. S. Nencka, Y. Wang, L. D. Nelson, M. A. McCrea, and T. B. Meier. Resting-state functional connectivity after concussion is associated with clinical recovery. *Hum. Brain Mapp.*, 40(4):1211–1220, Mar. 2019.
- [19] J. Kawahara, C. J. Brown, S. P. Miller, B. G. Booth, V. Chau, R. E. Grunau, J. G. Zwicker, and G. Hamarneh. BrainNetCNN: Convolutional neural networks for brain networks; towards predicting neurodevelopment. *Neuroimage*, 146:1038–1049, Feb. 2017.
- [20] J. Kim, J. R. Wozniak, B. A. Mueller, and W. Pan. Testing group differences in brain functional connectivity: using correlations or partial correlations? *Brain Connect.*, 5(4):214–231, May 2015.
- [21] J. D. Kruschwitz, D. List, L. Waller, M. Rubinov, and H. Walter. GraphVar: A user-friendly toolbox for comprehensive graph analyses of functional brain connectivity. *J. Neurosci. Methods*, 245:107–115, Apr. 2015.
- [22] J. P. Lachaux, E. Rodriguez, J. Martinerie, and F. J. Varela. Measuring phase synchrony in brain signals. *Hum. Brain Mapp.*, 8(4):194–208, 1999.
- [23] T. O. Laumann, A. Z. Snyder, A. Mitra, E. M. Gordon, C. Gratton, B. Adeyemo, A. W. Gilmore, S. M. Nelson, J. J. Berg, D. J. Greene, J. E. McCarthy, E. Tagliazucchi, H. Laufs, B. L. Schlaggar, N. U. F. Dosenbach, and S. E. Petersen. On the stability of BOLD fMRI correlations. *Cereb. Cortex*, 27(10):4719–4732, Oct. 2017.
- [24] G. Li, D. Han, C. Wang, W. Hu, V. D. Calhoun, and Y.-P. Wang. Application of deep canonically correlated sparse autoencoder for the classification of schizophrenia. *Comput. Methods Programs Biomed.*, 183(105073):105073, Jan. 2020.
- [25] H. Li, T. D. Satterthwaite, and Y. Fan. Brain age prediction based on resting-state functional connectivity patterns using convolutional neural networks. volume 2018 of *Proc. IEEE Int Symp Biomed Imaging*, pages 101–104, April 2018.
- [26] F.-H. Lin, K. Hara, V. Solo, M. Vangel, J. W. Belliveau, S. M. Stufflebeam, and M. S. Hämmäläinen. Dynamic Granger-Geweke causality modeling with application to interictal spike propagation. *Hum. Brain Mapp.*, 30(6):1877–1886, June 2009.
- [27] J. Liu, J. Ji, G. Xun, L. Yao, M. Huai, and A. Zhang. EC-GAN: Inferring brain effective connectivity via generative adversarial networks. *Proc. Conf. AAAI Artif. Intell.*, 34(04):4852–4859, Apr. 2020.
- [28] Q. Luo, W. Liu, L. Jin, C. Chang, and Z. Peng. Classification of obsessive-compulsive disorder using distance correlation on resting-state functional MRI images. *Front. Neuroinform.*, 15: 676491, Oct. 2021.
- [29] A. S. Mahadevan, U. A. Tooley, M. A. Bertolero, A. P. Mackey, and D. S. Bassett. Evaluating the sensitivity of functional connectivity measures to motion artifact in resting-state fMRI data. *Neuroimage*, 241(118408):118408, Nov. 2021.
- [30] P. R. Millar, P. H. Luckett, B. A. Gordon, and T. L. Benzinger. Predicting brain age from functional connectivity in symptomatic and preclinical alzheimer disease. *NeuroImage*, 256: 119228, 2022. ISSN 1053-8119. doi: <https://doi.org/10.1016/j.neuroimage.2022.119228>. URL <https://www.sciencedirect.com/science/article/pii/S1053811922003524>.

- [31] E. Mishor, D. Amir, T. Weiss, D. Honigstein, A. Weissbrod, E. Livne, L. Gorodisky, S. Karagach, A. Ravia, K. Snitz, D. Karawani, R. Zirler, R. Weissgross, T. Soroka, Y. Endevelt-Shapira, S. Agron, L. Rozenkrantz, N. Reshef, E. Furman-Haran, H. Breer, J. Strotmann, T. Uebi, M. Ozaki, and N. Sobel. Sniffing the human body volatile hexadecanal blocks aggression in men but triggers aggression in women. *Sci. Adv.*, 7(47):eabg1530, Nov. 2021.
- [32] M. M. Monti. Statistical analysis of fMRI time-series: A critical review of the GLM approach. *Front. Hum. Neurosci.*, 5:28, Mar. 2011.
- [33] A. Orlichenko, G. Daly, J. W. Freeman, and Y.-P. Wang. ImageNomer: developing an interactive graphical analysis tool for examining fMRI and omics data. In B. S. Gimi and A. Krol, editors, *Medical Imaging 2023: Biomedical Applications in Molecular, Structural, and Functional Imaging*. SPIE, Apr. 2023.
- [34] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Z. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. *CoRR*, abs/1912.01703, 2019. URL <http://arxiv.org/abs/1912.01703>.
- [35] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [36] U. Pervaiz, D. Vidaurre, M. W. Woolrich, and S. M. Smith. Optimising network modelling methods for fMRI. *Neuroimage*, 211(116604):116604, May 2020.
- [37] J. D. Power, A. L. Cohen, S. M. Nelson, G. S. Wig, K. A. Barnes, J. A. Church, A. C. Vogel, T. O. Laumann, F. M. Miezin, B. L. Schlaggar, and S. E. Petersen. Functional network organization of the human brain. *Neuron*, 72:665–678, 2011.
- [38] J. D. Ragland, B. I. Turetsky, R. C. Gur, F. Gunning-Dixon, T. Turner, L. Schroeder, R. Chan, and R. E. Gur. Working memory for complex figures: an fMRI comparison of letter and fractal n-back tasks. *Neuropsychology*, 16(3):370–379, July 2002.
- [39] D. Rangaprakash, R. Tadayonnejad, G. Deshpande, J. O’Neill, and J. D. Feusner. FMRI hemodynamic response function (HRF) as a novel marker of brain function: applications for understanding obsessive-compulsive disorder pathology and treatment response. *Brain Imaging Behav.*, 15(3):1622–1640, June 2021.
- [40] B. Rashid, M. R. Arbabshirani, E. Damaraju, M. S. Cetin, R. Miller, G. D. Pearlson, and V. D. Calhoun. Classification of schizophrenia and bipolar patients using static and dynamic resting-state fMRI brain connectivity. *Neuroimage*, 134:645–657, July 2016.
- [41] S. Sadeghi, D. Mier, M. F. Gerchen, S. N. L. Schmidt, and J. Hass. Dynamic causal modeling for fMRI with Wilson-Cowan-based neuronal equations. *Front. Neurosci.*, 14:593867, Nov. 2020.
- [42] T. D. Satterthwaite, M. A. Elliott, K. Ruparel, J. Loughhead, K. Prabhakaran, M. E. Calkins, R. Hopson, C. Jackson, J. Keefe, M. Riley, F. D. Mentch, P. M. A. Sleiman, R. Verma, C. Davatzikos, H. Hakonarson, R. C. Gur, and R. E. Gur. Neuroimaging of the philadelphia neurodevelopmental cohort. *NeuroImage*, 86:544–553, 2014.
- [43] K. E. Stephan and K. J. Friston. Analyzing effective connectivity with functional magnetic resonance imaging. *Wiley Interdiscip. Rev. Cogn. Sci.*, 1(3):446–459, May 2010.
- [44] D. O. Svaldi, J. Goñi, K. Abbas, E. Amico, D. G. Clark, C. Muralidharan, M. Dzemidzic, J. D. West, S. L. Risacher, A. J. Saykin, and L. G. Apostolova. Optimizing differential identifiability improves connectome predictive modeling of cognitive deficits from functional connectivity in alzheimer’s disease. *Hum. Brain Mapp.*, 42(11):3500–3516, Aug. 2021.
- [45] D. Szucs and J. P. Ioannidis. Sample size evolution in neuroimaging research: An evaluation of highly-cited studies (1990-2012) and of latest practices (2017-2018) in high-impact journals. *Neuroimage*, 221(117164):117164, Nov. 2020.
- [46] C. A. Tamminga, G. Pearlson, M. Keshavan, J. Sweeney, B. Clementz, and G. Thaker. Bipolar and schizophrenia network for intermediate phenotypes: outcomes across the psychosis continuum. *Schizophr. Bull.*, 40 Suppl 2:S131–7, Mar. 2014.

- [47] Y.-F. Tan, C.-M. Ting, F. Noman, R. C.-W. Phan, and H. Ombao. Graph-regularized manifold-aware conditional wasserstein GAN for brain functional connectivity generation. 2022.
- [48] B. O. Turner, E. J. Paul, M. B. Miller, and A. K. Barbey. Small sample sizes reduce the replicability of task-based fMRI studies. *Commun. Biol.*, 1(1):62, June 2018.
- [49] W. Yan, G. Qu, W. Hu, A. Abrol, B. Cai, C. Qiao, S. M. Plis, Y.-P. Wang, J. Sui, and V. D. Calhoun. Deep learning in neuroimaging: Promises and challenges. *IEEE Signal Processing Magazine*, 39(2):87–98, 2022. doi: 10.1109/MSP.2021.3128348.
- [50] X. Zhang, M. Liang, W. Qin, B. Wan, C. Yu, and D. Ming. Gender differences are encoded differently in the structure and function of the human brain revealed by multimodal mri. *Frontiers in Human Neuroscience*, 14, 2020.
- [51] J. Zhao, J. Huang, D. Zhi, W. Yan, X. Ma, X. Yang, X. Li, Q. Ke, T. Jiang, V. D. Calhoun, and J. Sui. Functional network connectivity (FNC)-based generative adversarial network (GAN) and its applications in classification of mental disorders. *J. Neurosci. Methods*, 341(108756): 108756, July 2020.

A Source Code

Code is available at <https://github.com/aorliche/AngleBasis>. Data used in our experiments is only available by request to the appropriate agencies, due to patient privacy issues (see Acknowledgements section in main paper). However, we provide sample preprocessed data taken from <https://openneuro.org/> in the GitHub repository. We also provide frozen models and Jupyter notebook code for the generation of synthetic FC based on user-provided subject characteristics.

B Effect of Jitter vs Angle on Decomposition

We investigated the possibility of using only jitter or only angle as a basis for decomposing FC. Although jitter could be estimated using gradient descent with constant phase angle, a more efficient way is to use the truncated eigendecomposition of the FC. For example, a rank-5 decomposition would be created as follows:

$$\begin{aligned}
 \mathbf{X} &= \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T \\
 \hat{\lambda}_{ii} &= \begin{cases} \lambda_{ii}, & i \leq 5 \\ 0, & \text{otherwise} \end{cases} \\
 \hat{\mathbf{X}} &= \mathbf{V}\hat{\mathbf{\Lambda}}\mathbf{V}^T
 \end{aligned} \tag{13}$$

where \mathbf{X} is a subject’s FC, $\hat{\mathbf{X}}$ is the low-rank estimate, and $\mathbf{\Lambda}$ is the ordered diagonal matrix of eigenvalues. This is possible because the FC is a positive semidefinite matrix.

Angle basis decomposition without jitter was estimated using gradient descent. The predictive accuracy of both models is shown in Figure 6.

C Comparison of Linear and Nonlinear Models for Downstream Tasks

We make a comparison between Logistic Regression, an MLP model, and the BrainNetCNN, similar to Pervais et al. [36]. As in that work, we find the linear model performed the best, with BrainNetCNN having worse performance (see Table 2). The small number of subjects and high inter-subject variability may hinder the advantage of more expressive models.

D Prediction using the Phase Lock Value Directly

We include in Figure 7 a set of representative PLV connectivities from the PNC dataset. As seen in Table 3, the PLV connectivities themselves are poor predictors of phenotypes.

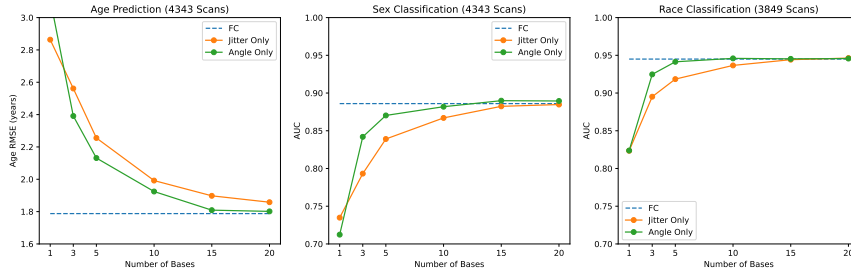


Figure 6: Prediction accuracy in the PNC dataset using only jitter vs using only the angle basis (cosine of phase difference).

Table 2: Comparison of Linear and Deep Models in BSNIP Dataset

Schizophrenia Diagnosis		
Model	AUC	p-value
Logistic Regression	0.785 ± 0.048	-
MLP	0.771 ± 0.025	0.128
BrainNetCNN	0.723 ± 0.028	< 0.001
Sex Prediction		
Model	AUC	p-value
Logistic Regression	0.755 ± 0.022	-
MLP	0.753 ± 0.012	0.90
BrainNetCNN	0.731 ± 0.049	0.069
Race Prediction		
Model	AUC	p-value
Logistic Regression	0.849 ± 0.011	-
MLP	0.849 ± 0.009	-
BrainNetCNN	0.755 ± 0.034	< 0.001

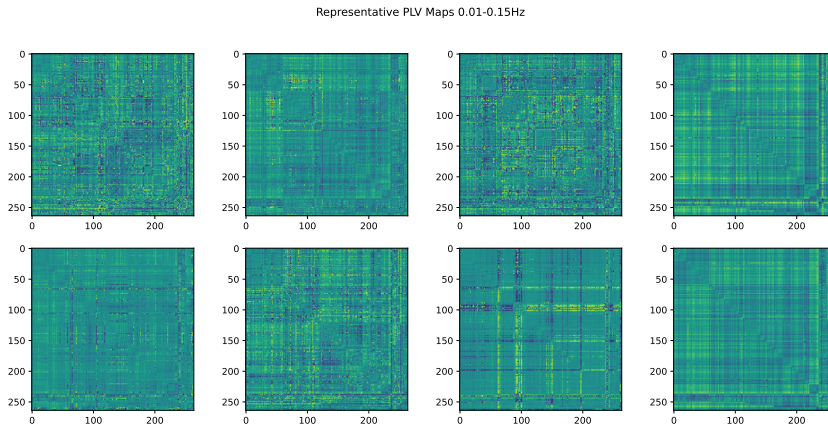


Figure 7: Phase lock value-based connectivities from the PNC dataset. Each image represents one subject.

Table 3: PLV Prediction Accuracy in the PNC Dataset

Metric	Predictive Task	FC	PLV	Null Model
RMSE (years)	Age	1.79 ± 0.02	3.13 ± 0.07	3.5
AUC	Sex	0.886 ± 0.006	0.605 ± 0.046	0.53
AUC	Race	0.946 ± 0.007	0.704 ± 0.033	0.52

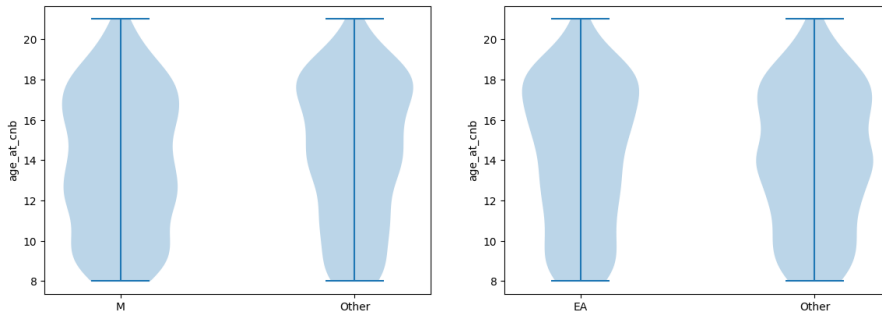


Figure 8: Sex and ethnicity broken down by age in the PNC dataset. EA is European Ancestry, Other is African Ancestry.

E Additional Demographic Information

We include more detailed demographic information on the PNC and BSNIP datasets in Figures 8, 9 and 10.

F Generation of Synthetic FC

Synthetic FC was generated by fitting 20 linear predictive models to the BSNIP dataset. The models used were Ridge Regression for age prediction and Logistic Regression for sex, race, and schizophrenia diagnosis prediction. Hyperparameters were the same as in the main paper. Models were fit on 20% random samples of the dataset.

One hundred synthetic subjects were created by initializing an angle basis model with 5 bases to random values, then performing gradient descent on the loss between the FC reconstructed from the angle basis and the specified values. A stop gradient operation was applied to model weights, so only the angle basis values (and consequently FC) could change. The phenotype parameters were set to 30 years for age and ±10 for the race, sex, and schizophrenia diagnosis logits. For reference, the absolute value of logits from the original predictive models ranged from 4 to 6 on their respective training sets. Gradient descent was performed for 100 epochs using the Adam optimizer with a

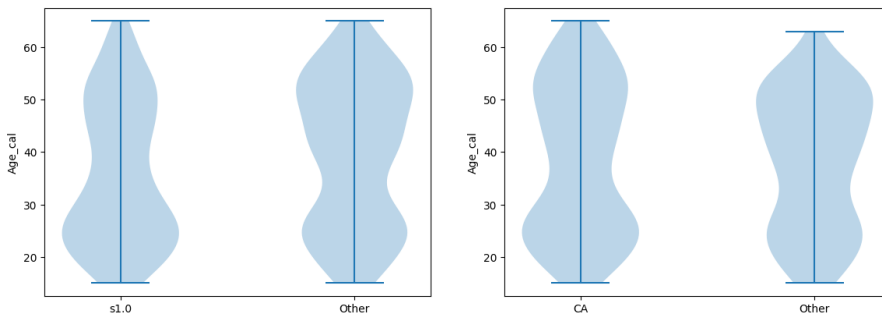


Figure 9: Sex and ethnicity broken down by age in the BSNIP dataset. “s1.0” refers to male and “Other” refers to female. CA is Caucasian Ancestry, Other is African Ancestry.

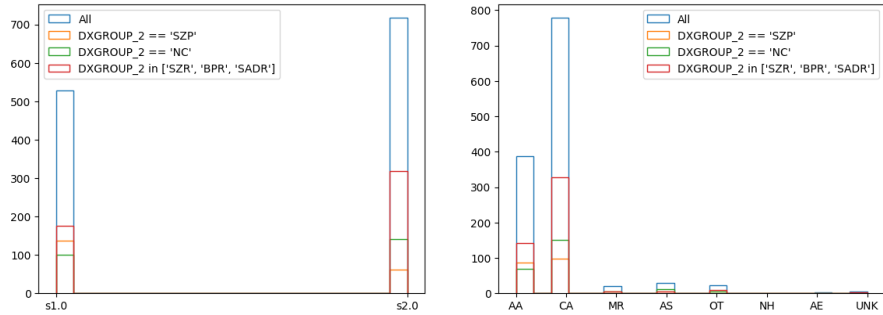


Figure 10: Sex and ethnicity broken down by patient subtype in the BSNIP dataset. “s1.0” refers to male, “s2.0” refers to female. Besides normal controls (NC) and patients (SZP, BPP, and SADR), there are also relatives of patients (SZR, BSPR, and SADR).

learning rate of 10^{-2} and no L2 regularization. The first 8 synthetic FCs for both groups are shown in Figure 11.

It should be noted that true subject FC exhibits far more variation than is captured in the synthetic data (see Figure 12). Nonetheless, we believe that the synthetic FC is useful since it accurately captures the coarse brain network-level differences between groups, and gives a simple visual description of qualitative differences in low-rank brain networks.

G Prediction with Limited Number of Features

We tested prediction on angle basis reconstruction and residual using features selected by Lasso, Ridge Regression, or correlation with the response variable. The results are shown in Figure 13. As expected, Lasso outperformed Ridge. Selecting features by correlation with the response variable was found to be a poor way to select features, likely due to the large multicollinearity present in FC.

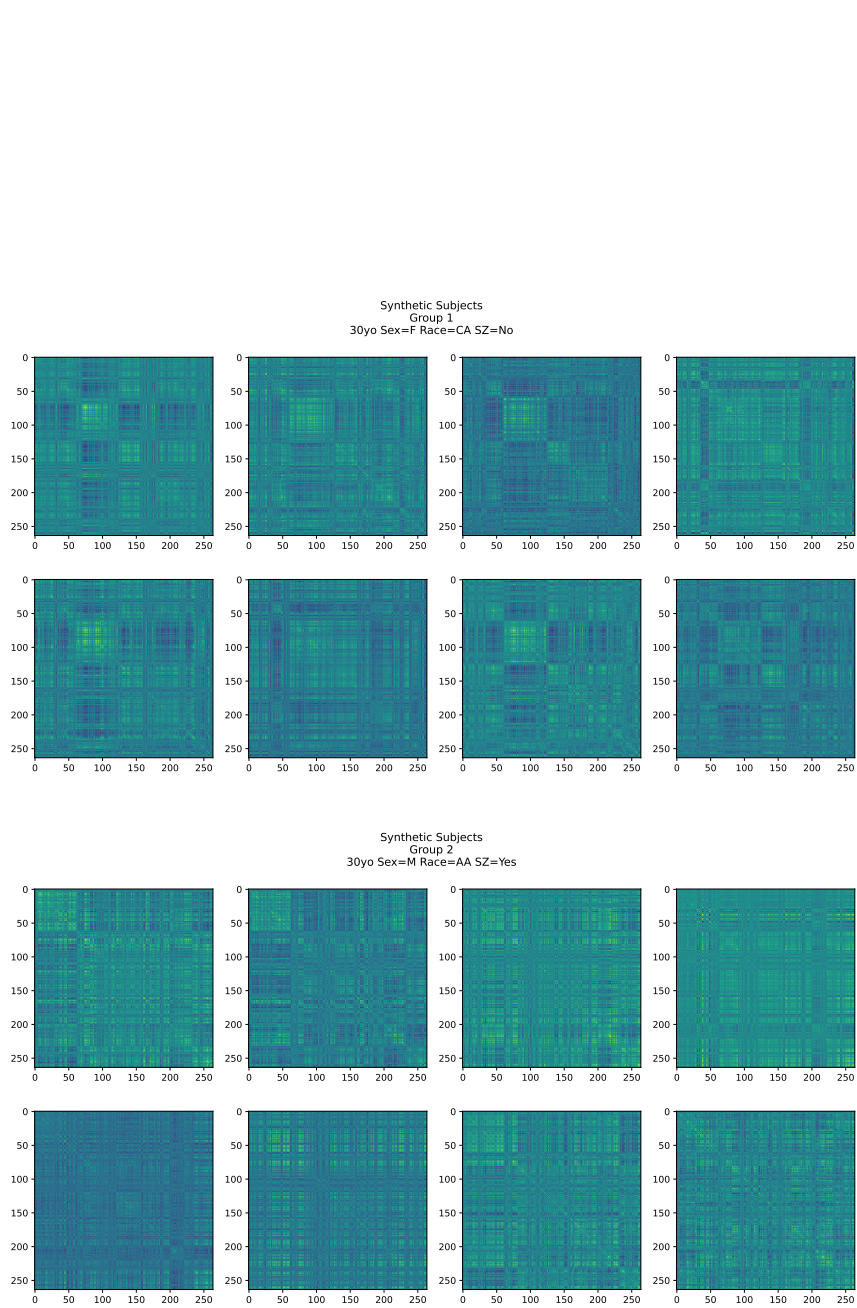


Figure 11: First 8 synthetic subjects from 30yo normal CA female group and 30yo schizophrenic AA male group.

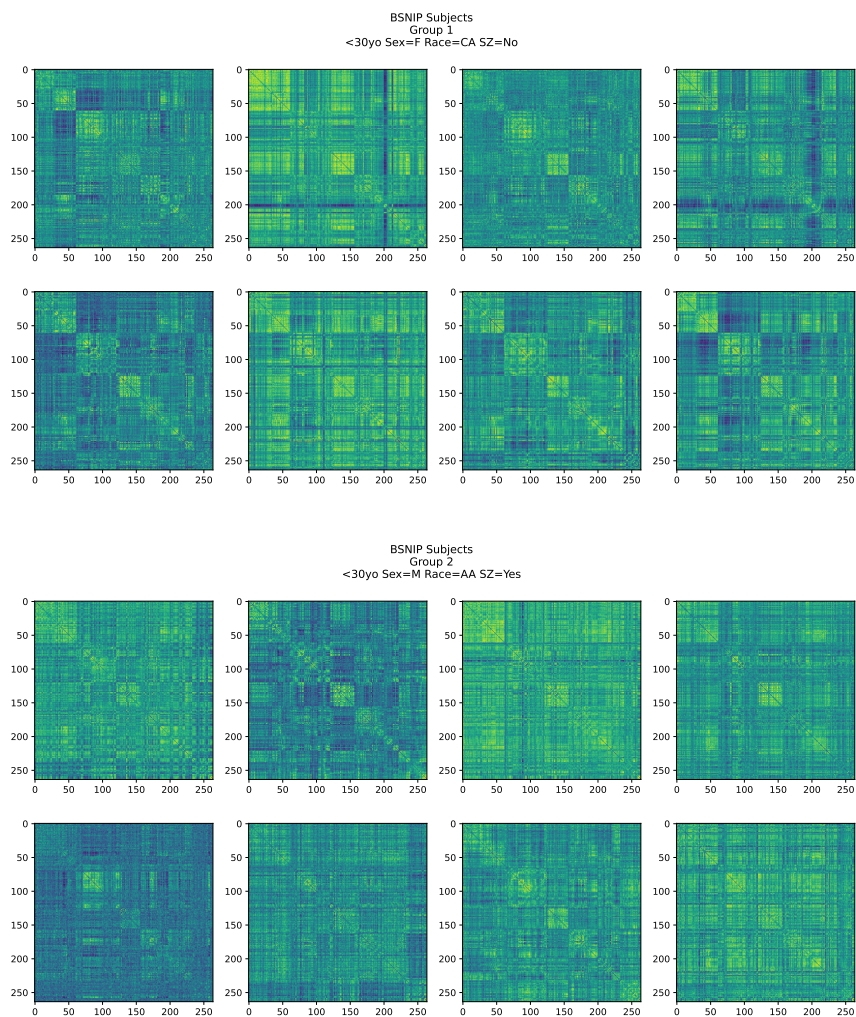


Figure 12: Example of the true variability of FC in the BSNIP dataset.

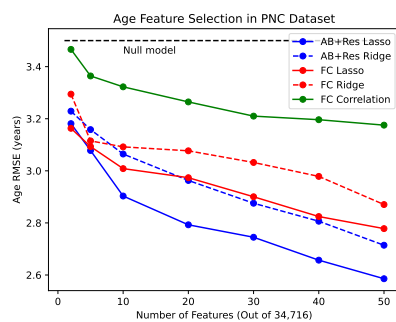


Figure 13: Feature selection for age prediction in the PNC dataset. Note features selected for angle basis reconstruction could be different than features selected for the angle basis residual.