Neural Networks 181 (2025) 106771

Contents lists available at ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

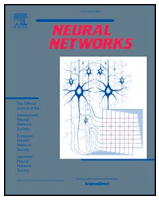

Classification of psychosis spectrum disorders using graph convolutional networks with structurally constrained functional connectomes

Madison Lewis [a], Wenlong Jiang [b], Nicholas D. Theis [c], Joshua Cape [d], Konasale M. Prasad [a,c,e,*]

[a] Department of Bioengineering, Swanson School of Engineering, University of Pittsburgh, Pittsburgh, PA 15213, United States
[b] Department of Statistics, University of Pittsburgh, Pittsburgh, PA 15213, United States
[c] Department of Psychiatry, University of Pittsburgh School of Medicine, Pittsburgh, PA 15213, United States
[d] Department of Statistics, University of Wisconsin–Madison, Madison, WI 52706, United States
[e] Veterans Affairs Pittsburgh Healthcare System, Pittsburgh, PA 15240, United States

ARTICLE INFO

Keywords:
Machine learning
Deep learning
Graph convolutional networks
Support vector machines
Psychotic-like experiences
Transdiagnostic sample

ABSTRACT

This article considers the problem of classifying individuals in a dataset of diverse psychosis spectrum conditions, including persons with subsyndromal psychotic-like experiences (PLEs) and healthy controls. This task is more challenging than the traditional problem of distinguishing patients with a diagnosed disorder from controls using brain network features, since the neurobiological differences between PLE individuals and healthy persons are less pronounced. Further, examining a transdiagnostic sample compared to controls is concordant with contemporary approaches to understanding the full spectrum of neurobiology of psychoses. We consider both support vector machines (SVMs) and graph convolutional networks (GCNs) for classification, with a variety of edge selection methods for processing the inputs. We also employ the MultiVERSE algorithm to generate network embeddings of the functional and structural networks for each subject, which are used as inputs for the SVMs. The best models among SVMs and GCNs yielded accuracies >63%. Investigation of network connectivity between persons with PLE and controls identified a region within the right inferior parietal cortex, called the PGi, as a central region for communication among modules (network hub). Class activation mapping revealed that the PLE group had salient regions in the dorsolateral prefrontal, orbital and polar frontal cortices, and the lateral temporal cortex, whereas the controls did not. Our study demonstrates the potential usefulness of deep learning methods to distinguish persons with subclinical psychosis and diagnosable disorders from controls. In the long term, this could help improve accuracy and reliability of clinical diagnoses, provide neurobiological bases for making diagnoses, and initiate early intervention strategies.

1. Introduction

Machine learning methods are increasingly being used in biomedical settings to improve diagnostic accuracy and reliability, prognostication, and to elucidate pathophysiology (Hassan et al., 2024; Mamdouh Farghaly et al., 2023; Yang et al., 2021). Inter-rater reliability and test-retest reliability of psychiatric diagnoses based on official classificatory systems such as the Diagnostic and Statistical Manual of Mental Disorders (DSM) are modest, and the biological validity of these diagnoses is unclear. Given the ability of machine learning methods to detect patterns in complex high-dimensional data, it is proposed that they may distinguish patients from controls with higher accuracy and reliability. Moreover, such prediction of diagnosis using network features adds biological validity to the clinical diagnostic approach that entirely depends on self-report of symptoms by patients and observations by experienced clinicians. The addition of biological validity may advance our understanding of the pathophysiology of psychotic disorders (Pereira et al., 2009).

This article considers the more challenging problem of classifying persons with psychotic-like experiences (PLE) and a transdiagnostic sample consisting of schizophrenia, major depressive disorder, and bipolar disorder. Persons with PLE have approximately 4 times higher risk for developing psychosis during the entire lifetime (Kaymaz et al., 2012) and 5–16-fold increase in risk among children (Welham et al., 2009). However, clinically distinguishing persons with PLE from healthy individuals is challenging, thereby highlighting the importance of

* Corresponding author at: 3811 O'Hara St, Pittsburgh, PA 15213, United States.
E-mail address: Kmp8@pitt.edu (K.M. Prasad).

https://doi.org/10.1016/j.neunet.2024.106771
Received 11 February 2024; Received in revised form 6 September 2024; Accepted 28 September 2024
Available online 30 September 2024

examining this population using machine learning methods to better distinguish PLE from controls. While the risk of conversion to psychosis is relatively low compared to the risk for psychosis in persons of clinical high risk (up to 35%) and familial high-risk (approximately 8%−15%), due to a higher prevalence of PLE in the community (7% of the general population (Linscott & van Os, 2013)), reliably identifying such persons is clinically significant as it may enable understanding early neurobiology before the onset of symptoms to help develop primary preventative strategies.

Using multimodal MRI data within machine learning models provides a non-invasive, data-driven, and neurobiology-based approach to classify persons with psychosis spectrum disorders (PSD). Clinically distinguishing PLEs from the general population is particularly challenging because most PLE individuals experience subclinical symptoms and do not seek treatment. In addition to higher prevalence in the community and higher conversion to psychosis (Linscott & van Os, 2013) as stated above, persons with PLE exhibit lower working memory capacity (Ziermans, 2013), higher frequency of anxiety and depression (Wigman et al., 2012), and decreased social functioning (Wikström et al., 2015). Using network data within machine learning methods can potentially lead to identifying persons with PLE and reveal early-stage pathophysiology without the confounds of medication exposure, since most individuals are not exposed to antipsychotics. Given that psychosis is considered a spectrum, ranging from mild subclinical symptoms to full-blown clinical syndrome, accurately classifying and examining network alterations can reveal the pathophysiological basis of psychosis spectrum at different severities of illness. This is important because psychotic disorders such as schizophrenia are chronic diseases with functional recovery as low as 14% (Jaaskelainen et al., 2013; Prasad, 2017; Robinson et al., 2004) and reduced life span of up to 25 years (Laursen et al., 2012).

In this article, we also examine a transdiagnostic sample consisting of schizophrenia, bipolar disorder, and major depressive disorder. These disorders are diagnostically and biologically overlapping, frequently difficult to clinically distinguish, and often are treated with the same medications. Transdiagnostic examination can identify common biological substrates that can distinguish these disorders from healthy control subjects.

Using data from the UK Biobank, we examined individuals with PLE who are included in the database as "unusual and psychotic experiences (UPE)," a transdiagnostic psychosis group consisting of schizophrenia, bipolar disorder, major depressive disorder defined according to the International Classification of Diseases, 10th edition (ICD-10), and a cohort of healthy subjects who were group-matched for age and sex with the latter two groups. We combined the PLE and the transdiagnostic groups to create a larger sample of psychosis spectrum disorder (PSD) subjects.

We implemented support vector machines (SVMs) and graph convolutional networks (GCNs) using inputs from structural MRI data (consisting of tractography data from the diffusion imaging) and resting-state functional MRI data, because prior studies show that both structural and functional networks are altered in persons with psychosis (Baker et al., 2014; Gong et al., 2017; Griffa et al., 2019; Rikandi et al., 2022).

SVMs are supervised machine learning models for discriminating between groups of data points and have been shown to classify major depressive disorder (Gallo et al., 2023), bipolar disorder (Achalia et al., 2020), and schizophrenia (Lei et al., 2022) with an accuracy of 61%, 87.6%, and 80.9%, respectively. The training phase locates the hyperplane that maximally separates the two groups, and a kernel can be applied to map the data into a higher-dimensional space to help identify the optimal discriminating hyperplane. Subsequently, a validation phase is used to adjust hyperparameters, and then a testing phase is used to measure classification accuracy. In this analysis, our SVM inputs were vectorized brain connectomes.

GCNs are an emerging state-of-the-art class of graph neural networks based on convolutional neural networks (CNNs) that were designed to aggregate local and neighborhood connection information to generate new feature maps. CNNs use filters on small regions of an image to identify local image features whereas GCN applies filters to each node's neighborhood (Khlifi et al., 2023). The convolutional step enables the algorithm to extract highly abstract features from the imaging data or graph connectivity networks but requires lengthy training and parameter tuning. Like SVMs, GCNs have been successful in classifying autism spectrum disorder (Ktena et al., 2018) and schizophrenia (Lei et al., 2022) from controls, as well as males and females using network features (Arslan et al., 2018). Unlike SVMs, GCNs consider neighborhood relationships from the graph structure instead of independent nodal edge weights. The GCN approach first uses an edge selection algorithm to eliminate unlikely connections in the graphs when forming the adjacency matrix. The adjacency matrix is then used as input for the convolutional layers, global average pooling layer, and a fully connected layer. A validation/testing set is used to measure model performance. Prior studies have employed only a single edge-selection method, whereas in this study, multiple edge-selection methods were employed. Namely, we considered methods which use functional information only, structural information only, and multiplex information (i.e., a combination of functional and structural), because edge selection methods can drastically alter the network structure that are inputted to the machine learning models in different ways. An additional advantage of GCNs is the ability to further differentiate the graph structure to identify which brain regions of the graph contributed the most to successful classification. This could in turn provide insight about which brain regions to target for novel treatment development.

The PLE networks may show less prominent changes compared to clinically diagnosed disorders, which have more prominent network alterations compared to healthy persons, because milder symptoms in PLE are difficult to clinically distinguish from healthy persons. Similarly, although the transdiagnostic disorders show dysfunctional activation in similar regions compared to controls, such dysfunctions are in the opposite direction for schizophrenia and bipolar disorder (McIntosh et al., 2008; Smucny et al., 2021; Theis et al., 2024). For these reasons, the dataset examined here is arguably more challenging to classify based on brain network features, so we did not necessarily expect the classification accuracy to be higher than prior studies classifying samples of clinically diagnosed subjects in comparison with healthy subjects. Further, machine learning methods could identify key brain regions/subnetworks that potentially discriminate the groups. We also hypothesized that GCNs would be more informative in terms of higher classification accuracy than SVMs, because GCNs utilize network structure with edge selection methods and can recognize features contributing to classification.

2. Methods

2.1. Data acquisition

Data was sourced from the UK Biobank (www.ukbiobank.ac.uk), a large-scale biomedical database containing lifestyle and health data on about half a million participants in the United Kingdom (Littlejohns et al., 2020), including MRI data for some participants (Miller et al., 2016). At the time we accessed the data for this study (March 2022), about 40,000 subjects had all three required scans for analysis: T_1-weighted, diffusion tensor imaging (DTI), and resting-state functional magnetic resonance imaging (rs-fMRI). Out of all subjects with required imaging modalities, we chose subjects experiencing PLE defined as UPE (data field 20,461; age when first experienced UPE) in the UK Biobank (n = 370) and subjects meeting ICD-10 diagnostic criteria for schizophrenia (data field 130,874), schizoaffective disorder (130,884), manic episodes (130,890), bipolar affective disorder (130,892), and/or recurrent depressive disorder (130,896). We grouped persons with these disorders together as the transdiagnostic group (n = 99). The

transdiagnostic group along with persons with lifetime experience of PLE were called psychosis spectrum disorder (PSD; $n = 469$). The control group ($n = 450$) comprised of individuals without PSD or PLE and was group-matched to PSD group for age and sex. The UK Biobank bulk data files were downloaded to the Pittsburgh Supercomputing Center (PSC) Bridges-2 system (Buitrago & Nystrom, 2021). Image acquisition parameters have been previously documented (Miller et al., 2016; Smith et al., 2022). Briefly, T_1w scans were acquired using 3D MPRAGE with voxel size=1 mm^3 and a TI/TR=880/2000 ms. The diffusion MRI had a voxel size of 2 mm^3 (isotropic) and a multiband factor (MB)=3. The TE/TR=39/735 ms, MB=8, flip angle=52° for resting-state fMRI acquisition.

2.2. Anatomical image processing

Specialized neuroimaging programs used for additional processing were available on the Pittsburgh Supercomputing Center (Bridges-2 system) using official Singularity containers (Kurtzer et al., 2017). The preprocessing pipeline for all modalities is presented in Fig. 1. The UK Biobank image files titled "T1_unbiased_brain.nii.gz" were used for the anatomical MRI (aMRI) processing for gradient distortion correction, brain extraction, and bias correction by the UK Biobank group (Smith et al., 2022). T_1-weighted images were processed on Bridges-2 as part of the multiplex pipeline using the recon-all function from FreeSurfer (version 7.2.0) via the official dockerized image, using Singularity Image Format (SIF) (Fischl, 2012). Then, each individual brain was parcellated according to the Human Connectome Project Multi-Modal Parcellation (HCP MMP1) atlas (Glasser et al., 2016), using a published method (Neurolab, 2017). Briefly, this method uses a Freesurfer (FS) average version of the HCP MMP1 parcellation (Mills, 2016) to the individual result of recon-all which produces the standard 360-node parcellation (180 parcels per hemisphere). Since two of the Glasser atlas regions are the left and right hippocampus and these are included in the Freesurfer-derived subcortical structures, we removed them to keep 358 *cortical* parcels. Next, we included the 21 subcortical parcellations which included bilateral hippocampi. Because two subcortical structures from the automated segmentation were too small to be reliably parcellated (the left and right substantia nigra), these were removed to yield 19 subcortical structures, resulting in a total of 377 parcels as network nodes.

2.3. DTI processing

DTI data was preprocessed by the UK Biobank (Smith et al., 2022) for eddy current correction, head motion correction, and gradient distortion correction. We performed fiber tracking using the Mrtrix3 (Tournier et al., 2019) official docker image https://hub.docker.com/r/mrtrix3/mrtrix3 via Singularity on Bridges-2. First, a brain mask was created from the preprocessed DTI. Next, a response function was calculated as well as a Fiber Orientation Distribution (FOD) image by spherical deconvolution of preprocessed DTI with the response function. Fiber tracking was then performed on the FOD image under the brain mask by drawing 10 million streamlines with a maximum of 1000 attempts each (maximum angle of 45°) using anatomically constrained tractography (ACT) by the five-tissue-type (5TT) model. Tracts were also reduced using Spherical-deconvolution Informed Filtering of Tractograms (SIFT).

The native space (T_1 space) brain was then registered to the DTI space image using FSL FLIRT linear interpolation. The resulting transform matrix used to register the aMRI to the DTI was then applied to the

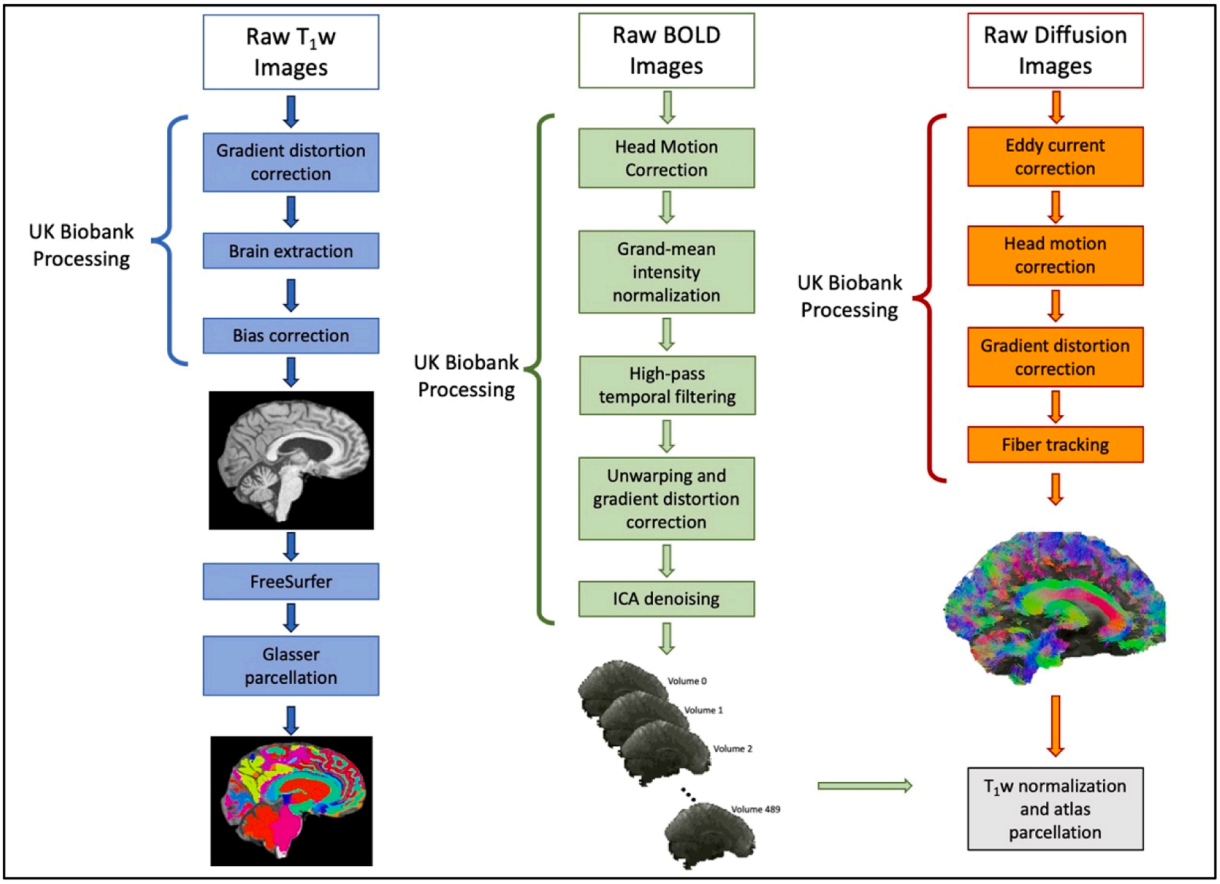

Fig. 1. Preprocessing pipeline for T1w images (left, blue), functional bold images (middle, green), and diffusion images (right, orange). Preprocessing steps completed by the UK Biobank are indicated using brackets.

atlas parcellation image using nearest-neighbor interpolation to provide a DTI-space anatomical atlas. This registration step was then quality controlled by visual inspection. Finally, the Mrtrix3 function "tck2connectome" was used to generate an adjacency matrix – a structural connectome (SC) – where edge weights represent streamline counts (after SIFT). All SCs were harmonized to eliminate multisite effects using ComBat (Yu et al., 2018).

2.4. fMRI image processing

Resting state fMRI images were preprocessed by the UK Biobank (Smith et al., 2022) for head motion (MCFLIRT) (Jenkinson et al., 2012), grand-mean intensity normalization of each 4D dataset, high-pass temporal filtering, echo planar imaging (EPI) unwarping and gradient distortion correction, and finally independent component analysis (ICA)-based denoising (Beckmann & Smith, 2004; Griffanti et al., 2014; Salimi-Khorshidi et al., 2014). The native T_1 brain space was registered to the rs-fMRI space image using FSL FLIRT linear interpolation. The resulting transform matrix used to register the aMRI to the rs-fMRI was then applied to the atlas parcellation image using nearest-neighbor interpolation to provide a rs-fMRI-space anatomical atlas. This registration step was then quality controlled by visual inspection. Finally, the spatially averaged time series (across all voxels in a parcel) was then calculated for each node, and the temporal correlations of nodal time series (Pearson's correlation, r) were then determined for each node pair, resulting in a resting state functional connectome (FC). All FCs were harmonized to eliminate multisite effects using ComBat (Yu et al., 2018).

2.5. Edge-Selection methods

Edge-selection methods, also known as thresholding methods, were employed to select network edges that represent highly probable connections and to remove edges that are most likely attributed to noise which therefore may not represent true functional edges. The choice of edge-selection method can significantly change the graph properties of the original network. To test which method is best for distinguishing functional connectivity between the groups, we compared multiple commonly used methods. The K-nearest neighbors (KNN) method (Cover & Hart, 1967) uses Euclidean distances between pairs of nodes to estimate similarity and has achieved a high success rate in prior studies using graph neural networks (Kang, 2021). We chose the values $k = 10$, 50, and 100 since our network size was significantly larger than in the prior studies.

In addition to KNN, we implemented the percolation threshold which determines the largest magnitude threshold at which the network's giant connected component (GCC) contains all nodes in the network (Bordier et al., 2017). This method aims to find the optimal balance between information gained by noise reduction and information lost by excessive pruning.

We also implemented a method developed by our group, namely objective function threshold method (Theis et al., 2023), which finds a threshold value that optimizes a given graph metric calculation. Our previous work demonstrated that characteristic path length is a useful choice of metric over which to optimize. The objective function threshold method can yield binary networks with one or more connected components, depending on the choice of parameter value.

Given an input network, Kruskal's Algorithm (Kruskal, 1956), as implemented in the Maximum Weight Spanning tree (Undirected) MATLAB package (Chu & Liu, 1965; Edmonds, 1967; Li, 2024) was used to determine the Maximum Spanning Tree (MST). The MST is defined as the smallest subset of edges in a graph that connects all graph nodes (GCC=1) such that the corresponding total edge weight is maximal. A simple and widely used thresholding method is proportional thresholding, implemented in the brain connectivity toolbox (BCT) (Rubinov & Sporns, 2010). This technique removes the weakest set of edges necessary to achieve a desired network density. We combined the MST approach and the proportional thresholding approach to retain edges with the largest weight after the MST edge to make the MST denser. We called this edge-selection method "density-matched Maximum Spanning Subgraph" (dMSS). The proportion of non-MST edges added back to the FC is determined by the density of the corresponding SC for the same subject.

Additionally, we included a multiplex method that combines both SC and FC information as inputs to the GCN model. We start with subjects' SC to define a binary network of all node-pairs connected by at least one diffusion streamline, thus yielding sparse networks with approximately 10% edge density. Since our input graph for each subject is the complete FC, we additionally supply this SC-derived binary matrix as the edge selection scheme. This forces the GCN model to only consider the FC connections that have a corresponding SC connection. Hereafter we refer to this method as "SC-constrained" edge selection.

2.6. SVMs and SVMs with multiverse

In the SVM analysis, for each subject the rs-FC and SC represented the weighted adjacency matrices and were used to build separate functional and structural classifiers. Namely, we vectorized the upper triangle elements of the adjacency matrix so that for each sample i, we have $n(n-1)/2$ features denoted as $\boldsymbol{X}_{FC}^{[i]}$ and $\boldsymbol{X}_{SC}^{[i]}$. We then train the SVMs using $\boldsymbol{X}_{FC}^{[i]}$ and $\boldsymbol{X}_{SC}^{[i]}$ separately. However, there are two challenges when using this method. First, the number of features is very large (377 nodes creates 70,876 possible features) compared to other studies that have a smaller network size. Second, we did not combine the information from the FC and SC to train the SVMs because this would double the number of features (which is already large), so instead we used the multiplex network embedding approach as described below.

To combine information from the FC and SC, we used the MultiVERSE algorithm (Pio-Lopez, Valdeolivas, Tichit, Remy, & Baudot, 2021) which builds an embedding matrix from multiplex SC and FC networks. The MultiVERSE algorithm involves two main steps. First, given a structural and functional network, we constructed the similarity measure between two nodes using a random walk with restart. Next, given the similarity measure, we obtained an embedding matrix $\boldsymbol{W} \in \mathbb{R}^{n \times d}$ where $d < n$ by using the noise contrastive estimation method (Gutmann & Hyvärinen, 2010) used within the MultiVERSE algorithm. For clarity, we followed the published procedure (Pio-Lopez, et al., 2021). For each sample i, we have a two-layer multiplex network where the first layer is the functional network, and the second layer is the structural network to find a matrix $\boldsymbol{W}^{[i]}$ and then vectorize it into the feature vector $\boldsymbol{X}^{[i]} \in \mathbb{R}^{nd}$ (Pio-Lopez, et al., 2021). We used unweighted FC and SC with equal graph density as our inputs to the MultiVERSE algorithm. For the SC, we considered an edge to be present when at least one streamline count was observed, whereas for the FC we used the dMSS thresholding described above. We used the Python code in Pio-Lopez et al. to implement MultiVERSE using the recommended default values for most parameters (Pio-Lopez, et al., 2021). We considered different embedding dimensions ($d \in 2, 35, 128$), where 128 is the default dimension used in the MultiVERSE algorithm.

The traditional SVM with linear kernel was used to construct a maximal margin hyperplane based on the training data, because previous classification studies using neuroimaging data with this choice reported good performance (Achalia, et al., 2020; Gallo, et al., 2023; Lei, et al., 2022). We also considered SVMs with the radial kernel function, which constructs a nonlinear boundary by enlarging the feature space. For all input types, FC only, SC only, and MultiVERSE, the same tuning parameters for SVMs were used. For the linear kernel, we have one tuning parameter C and for the radial basis function (RBF) kernel, we have two tuning parameters C and γ. We considered $C \in \{0.0001, 0.001, 0.01, 0.1, 1, 10, 100, 1000\}$ and $\gamma \in \{0.001, 0.01, 0.1, 1, 10, 100\}$. To evaluate SVM performance, pooled stratified 10-fold

cross-validation was used to split the data into training/validation and test sets. The training set was used to train the SVMs, the validation set was used to select the tuning parameters, and the test set was used to evaluate the classification performance of the SVMs (see Supplemental table 1 for sizes of the training, validation, and test sets).

2.7. Graph convolutional networks

We modified an existing GCN pipeline (Lei et al., 2022) to examine single subject classification of controls and PSD (Fig. 2). Modifications were made to the type of edge selection method used, input network architecture, and therefore how the data was loaded into Pytorch. The GCN model was kept the same except for the first input to convolutional layer 1 having a larger number of features. In short, three graph convolutional layers (number of channels: [64, 64, 128]) were used in our model followed by a global average pooling layer and a fully connected layer. The Rectified Linear Unit (ReLU) activation function was used to add hidden layers and introduce non-linearity into the model. The Softmax function encoded the output value into a predictive probability for each class from the output layer. In our analysis, the FC was represented using graph structure where each region is a node, and each node feature is the functional connectivity of that region and all other regions. The adjacency matrix was calculated via the edge-selection methods discussed above. For completeness, the SC consisted of the same nodes as the FC with each node feature representing the streamline count of that region to all other regions. The adjacency matrix was calculated using a threshold of streamline counts ($\geq$1). To evaluate GCN performance, pooled stratified 10-fold cross-validation was used to split the data into training and test sets. For each fold, 50 epochs of the training dataset were run through the learning process. The performance of each epoch was determined by calculating training loss, validation loss, balanced accuracy, sensitivity, and specificity, and test balanced accuracy, sensitivity, and specificity (see Appendix 1 for the GCN pseudocode). For more information on network architecture, see pseudocode part 2 of the Appendix 1; see Fig. 2D for information on convolutional functions and feature dimensions, or Supplemental Fig. 1 for trainable parameters. Sizes of the training, validation, and test sets are described in Supplemental Table 1.

2.8. Post HOC analysis for PLE and transdiagnostic samples

After examining the entire PSD group compared with healthy controls, we conducted a post hoc investigation of the transdiagnostic sample and the PLE group separately, compared to age and sex matched controls. The balanced accuracy was calculated in the same way as above. Since the sample size of the transdiagnostic group was smaller than the PSD and PLE groups, we chose to use $k = 5$ instead of 10 for the k-fold cross validation. The connectome and edge selection method which performed the best on the PSD sample was used for the GCN inputs and adjacency matrix calculation for PLE and transdiagnostic groups (see Supplemental Table 1 for sizes of the training, validation, and test sets).

2.9. Network analysis for graph theoretic features

To further understand network differences between the patient and control groups, the model with the highest accuracy was used to classify a random set of subjects that was the same size as the test set. The subjects were then split into groups based on whether they were classified correctly or incorrectly as patient or control, resulting in four groups: true-patient, true-control, false-patient, and false-control. These four groups were analyzed separately using graph theoretic methods.

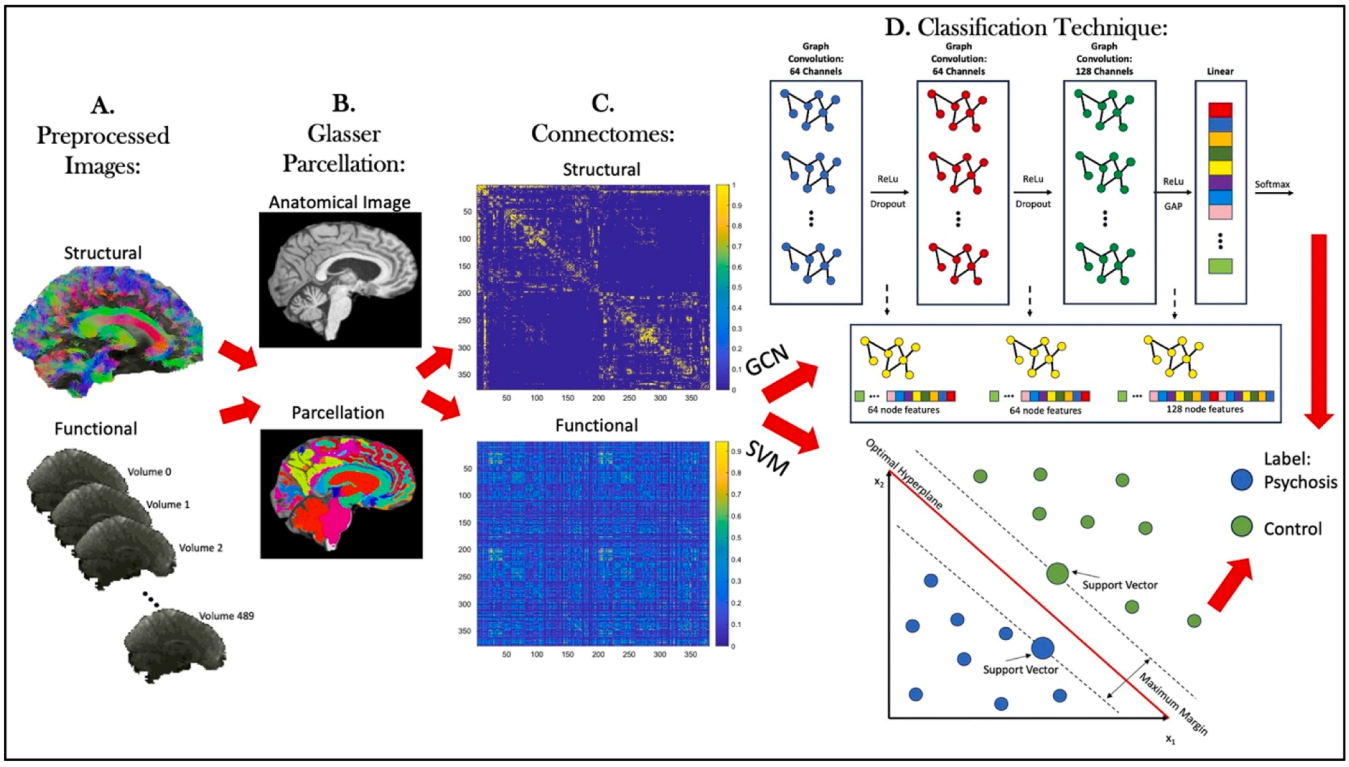

Fig. 2. Pipeline used to classify preprocessed structural and functional images. (A) First the anatomical, diffusion, and functional images are preprocessed (see also Fig. 1). The preprocessing outputs of the diffusion, a tractogram, and the functional, a denoised and realigned timeseries, data are shown. (B) The Glasser atlas (cortical) and FreeSurfer (subcortical) were used to parcellate the brain (T_1w image, top) into 377 regions of interest. The parcellation was then applied to the outputs of part A. (C): Connectomes were built using streamline count (diffusion) and Pearson correlation (functional) between each region and all other 376 regions. This step creates a 377-by-377 connectivity network from diffusion and functional data separately. (D) After network creation, Support Vector Machines (SVM) and Graph Convolutional Networks (GCN) were used to classify psychosis from controls.

For each subject in each group, network hubs were calculated and the inter-section connections, based on the Glasser atlas section definitions, were compared. Hubs were defined as nodes with degree, betweenness centrality, or eigenvector centrality more than two standard deviations above the network average of the same measure (Bassett et al., 2008; Lewis et al., 2023). Global graph centrality (average eigenvector and betweenness centrality) and segregation (modularity) measures were calculated. These were then compared between groups to quantify differences in segregation and centrality metrics. Of note, modularity relates the number of within-network connections to all connections to quantify the strength of segregation into distinct networks (Cohen & D'Esposito, 2016). Average eigenvector and betweenness centrality were used to measure a network's regional interaction and resilience (Rubinov & Sporns, 2010).

2.10. Salient region detection using class activation mapping (CAM)

In addition to investigating graph theoretic properties, we used Class Activation Mapping (CAM) to detect salient regions that were informative for classification (Arslan et al., 2018). Using the GCN model with the highest accuracy, we extracted the final dense layer, or the output from the final convolutional layer, from all subjects. To estimate the activation value of each node, we multiplied each subject's final layer (in our case 377 × 128 matrix) by the optimized weight vector for their respective class (128 × 1 vector). This yielded a 377 × 1 CAM vector for each subject in each class. Here, we were interested in analyzing the two group's activation values separately, so we calculated the average CAM vector across subjects in each group. The top 10 nodes with the highest activation value are reported from each group.

3. Results

3.1. Demographic data

After extensive processing and quality control of processed outputs as detailed above, a final sample of 919 multiplex networks were used (PSD=469; controls=450). Of the 469 PSD subjects, 370 had lifetime experience of PLE but were never diagnosed with an ICD disorder (PLE sample), and 99 had an ICD-10 disorder (transdiagnostic sample). Control and PSD did not significantly differ by age. The sex distribution between the PSD (72% female) and the control group (64% female) and between the PLE and controls were significantly different, respectively, but not between the transdiagnostic and controls (Table 1). The same training, validation, and test sets were used for both GCNs and SVMs for k-fold cross validation (Supplemental Table 1).

3.2. SVMs and SVMs with MultiVERSE

The SVMs correctly classified the PSD group with 56% average accuracy across the folds that used FC edge weights. The models using SC and MultiVERSE embedding performed relatively poorly (Table 2A). However, the best model using the FC edge-weight inputs had an accuracy of 65% (Table 2B). The linear kernel consistently outperformed the radial basis function (RBF) kernel (Supplemental Table 2).

Table 2
PSD classification accuracy, sensitivity, and specificity for the test sets using SVM. (A) Average metrics across all 10 folds for functional, structural, and MultiVERSE inputs. (B) Metrics for the best model, the model with the highest accuracy, for the three different inputs. The highest average and best model accuracies are in bold. Metrics using a linear kernel are shown for all SVM models since it outperformed RBF. The best embedding dimension for MultiVERSE is shown. Metrics for other embedding dimensions and the RBF kernel is provided in the Supplemental Table 2.

A. SVM Average Metrics Across Folds for Psychosis Spectrum Disorder Sample

Input	Average Accuracy	Average Sensitivity	Average Specificity
Functional Connectome	**0.56±0.06**	0.56±0.06	0.55±0.10
Structural Connectome	0.50±0.05	0.51±0.06	0.49±0.08
MultiVERSE Embedding ($d = 16$)	0.51±0.06	0.50±0.08	0.52±0.07

B. Best SVM Model Metrics for Psychosis Spectrum Disorder Sample

Input	Accuracy	Sensitivity	Specificity
Functional Connectome	**0.65**	0.62	0.69
Structural Connectome	0.59	0.62	0.56
MultiVERSE Embedding ($d = 2$)	0.61	0.70	0.51

3.3. Graph convolutional network models

Across the folds, SC input provided the highest average accuracy (53%) with 38% sensitivity and 68% specificity (Table 3A). Balanced accuracy varied marginally (~2%) based on the connectome input and edge selection method chosen, but some provided imbalanced sensitivity and specificity. Similarly, the FC input with SC-constrained edge selection yielded 52.5% accuracy with 47% sensitivity and 58% specificity (Table 3A). In terms of best model metrics, the lowest balanced accuracy (57.9%) was achieved using the KNN algorithm, and the best classification accuracy was achieved using SC-constrained (62.0%) and had a balanced sensitivity and specificity (Table 3B). KNN with $k = 100$ performed the best but not $k = 10$ or 50 that performed <50% accuracy and are therefore not reported.

3.4. Post HOC analysis for PLE and transdiagnostic samples

We examined the diagnostic components of the PSD group, namely the PLE and the transdiagnostic samples separately. The highest average accuracy for the PLE across folds was 54% and was achieved using linear kernel SVMs with functional connectome edge weights as the input (Table 4A). The best model, with an accuracy of 65%, was MultiVERSE SVM using an embedding dimension=8 with a linear kernel (Table 4B). See Supplemental Table 3 for metrics for the other embedding dimensions and the RBF kernel.

For the transdiagnostic sample, the highest average accuracy (57%) using MultiVERSE with SVM was achieved using the embedding dimension $d = 4$ and a linear kernel (Table 5A). The best model across methods and folds was also using MultiVERSE embedding with a dimension $d = 4$ as the input and had a 69% accuracy (Table 5B; see Supplemental Table 4 for metrics for other embedding dimensions and the RBF kernel).

Since SC and FC with SC-constrained edge selection inputs provided

Table 1
Demographic data of the samples used for classification tasks.

	PSD Sample		PLE Sample		Transdiagnostic Sample	
Group	PSD ($n = 469$)	Control ($n = 450$)	PLE ($n = 370$)	Control ($n = 370$)	Transdiagnostic ($n = 99$)	Control ($n = 99$)
Age (years)	53.2 ± 6.9	53.6 ± 7.3	53.3 ± 6.7	53.3 ± 7.4	52.6 ± 7.5	53.0 ± 6.8
Age Comparison	$t = 0.98$	$p = 0.91$	t=−0.12	$p = 0.91$	$t = 0.38$	$p = 0.71$
Sex (M/F)	131/338	162/288	101/269	131/239	31/68	40/59
Sex Comparison	χ^2=6.76	$p = 0.0093$	χ^2=5.64	$p = 0.0175$	χ^2=1.63	$p = 0.201$

Table 3
Graph Convolutional Network metrics for PSD sample. (A) Average balanced accuracy, sensitivity, and specificity across the 10-folds for different edge selection methods. (B) Metrics for the model with best performance, based on highest accuracy. The highest average and best model accuracies are in bold. KNN= *k*-nearest neighbors ($k = 100$). For more information on each edge selection, see the methods section.

A. GCN Average Metrics Across Folds for Psychosis Spectrum Disorder Sample

Connectome	Edge-Selection	*Av.* Accuracy	*Av.* Sensitivity	*Av.* Specificity
Functional	KNN	0.52±0.04	0.38±0.27	0.66±0.24
Functional	Objective Function	0.51±0.05	0.34±0.16	0.68±0.19
Functional	Percolation	0.52±0.04	0.60±0.17	0.43±0.15
Multiplex (structural + functional)	dMSS	0.49±0.05	0.39±0.27	0.59±0.31
Multiplex (structural + functional)	SC-constrained	0.53±0.07	0.47±0.25	0.58±0.27
Structural	Streamline Count≥1	**0.53 ±0.05**	0.38±0.15	0.68±0.12

B. Best GCN Model Metrics for Psychosis Spectrum Disorder Sample

Connectome	Edge-Selection	Accuracy	Sensitivity	Specificity
Functional	KNN	0.58	0.87	0.29
Functional	Objective Function	0.61	0.75	0.47
Functional	Percolation	0.61	0.60	0.62
Multiplex (structural + functional)	dMSS	0.58	0.40	0.76
Multiplex (structural + functional)	SC-constrained	**0.62**	0.62	0.62
Structural	Streamline Count≥1	0.61	0.53	0.69

Table 4
PLE classification metrics for the test sets using SVMs. (A) Average metrics across all 10 folds for functional, structural, and MultiVERSE inputs. (B) Metrics for the best model, the model with the highest accuracy, for the three different inputs. The highest average and best model accuracies are in bold. Metrics using a linear kernel are shown for all SVM models since it outperformed RBF. The best embedding dimension for MultiVERSE is shown. Metrics for other embedding dimensions and the RBF kernel can be viewed in the supplemental material.

A. SVM Average Metrics Across Folds for PLE Sample

Input	Average Accuracy	Average Sensitivity	Average Specificity
Functional Connectome	**0.54 ± 0.06**	0.49 ± 0.16	0.58 ± 0.11
Structural Connectome	0.49 ± 0.06	0.50 ± 0.12	0.47 ± 0.07
MultiVERSE Embedding ($d = 2$)	0.53 ± 0.03	0.52 ± 0.09	0.53 ± 0.08

B. Best SVM Model Metrics for PLE Sample

Input	Accuracy	Sensitivity	Specificity
Functional Connectome	0.62	0.54	0.70
Structural Connectome	0.58	0.68	0.49
MultiVERSE Embedding ($d = 8$)	**0.65**	0.62	0.68

the highest accuracy for the PSD samples, these two edge selection methods were used for classifying the PLE and transdiagnostic groups. The balanced accuracy improved marginally when GCN was implemented on these subgroups. For the PLE-only sample, the best average balanced accuracy was 54%, which was similar to the PSD, but the best model accuracy showed a marginal improvement at 65% compared to 62% (Table 6). Similarly, when the transdiagnostic sample was examined, the best average balanced accuracy (55%) improved by 2% and 1% for the best model (Table 6).

Table 5
Transdiagnostic sample classification metrics for the test sets using SVMs. (A) Average metrics across all 5 folds for functional, structural, and MultiVERSE inputs. (B) Metrics for the best model, the model with the highest accuracy, for the three different inputs. The highest average and best model accuracies are in bold. Metrics using a linear kernel are shown for all SVM models since it outperformed RBF. The best embedding dimension for MultiVERSE is shown. Metrics for other embedding dimensions and the RBF kernel can be viewed in the supplemental material.

A. SVM Average Metrics Across Folds for Transdiagnostic Sample

Input	Average Accuracy	Average Sensitivity	Average Specificity
Functional Connectome	0.51 ± 0.08	0.69 ± 0.31	0.33 ± 0.39
Structural Connectome	0.55 ± 0.09	0.57 ± 0.09	0.53 ± 0.11
MultiVERSE Embedding ($d = 4$)	0.57 ± 0.10	0.56 ± 0.14	0.59 ± 0.08

B. Best SVM Model Metrics for Transdiagnostic Sample

Input	Accuracy	Sensitivity	Specificity
Functional Connectome	0.62	0.30	0.95
Structural Connectome	0.69	0.68	0.70
MultiVERSE Embedding ($d = 4$)	0.69	0.70	0.68

3.5. Network analysis for graph theoretic features

GCNs yielded the best model (65% accuracy) using the SC with the streamline count≥1 as the edges to classify persons with PLE from controls. This model's test subjects' thresholded networks were extracted and analyzed for differences in graph measures. None of the global graph measures were significantly different between PLE and control groups. The right PGi, part of the inferior parietal cortex, was identified as a hub in PLE subjects but not in controls. The right PGi was considered to be a connector hub since it had first-degree connections with other inferior parietal regions, lateral temporal cortex, temporo-parieto-occipital junction, and auditory association cortex (Fig. 3) that are different cortical sections, defined by the Glasser parcellation, allowing for communication between these sections (Farahani et al., 2019). There were 10 hubs in all groups that did not impact model classification (Supplementary Table 5; see Supplemental Materials for results from the PSD and transdiagnostic samples).

3.6. Salient region detection using class activation mapping (CAM)

Using the GCN model for PLE that yielded the highest accuracy (65%), we extracted each subject-specific CAM vector. Due to the relatively low classification accuracy, sensitivity, and specificity, we examined the groups separately. For the PLE and control groups, the top 10 most salient regions were located in different cortical sections. Both groups had salient regions in the primary visual, early visual, somatosensory, and motor, and subcortical Glasser-defined sections. The PLE group had salient regions in dorsolateral prefrontal cortex, lateral temporal, and orbital and polar frontal cortex but the control group had salient regions in auditory association, posterior cingulate, inferior frontal, paracentral lobule and mid cingulate cortex, and superior parietal cortex sections (Fig. 4). Results from the PSD and transdiagnostic samples are available in Supplementary Material.

4. Discussion

Our main findings are that SVMs and GCNs provided similar prediction accuracy using network characteristics for PSD (PLE + transdiagnostic group). The best model for SVMs achieved 65% classification accuracy for the PLE and 69% for the transdiagnostic sample, whereas GCNs achieved 65% accuracy for the PLE and 63% for the transdiagnostic sample. That being said, the average accuracy across all folds for SVM and GCN were similar. The PLE group showed an accuracy

Table 6
GCN performance after dividing the PSD sample into more distinct subgroups: PLE and transdiagnostic samples. (A) Average balanced accuracy, sensitivity, and specificity across the 10-folds (PLE) and 5-folds (transdiagnostic) for different edge selection methods. (B) Metrics for the model with best performance, based on highest accuracy. The highest average and best model accuracies are in bold. The most distinct group is the transdiagnostic sample (gray filled) and the PLE group (white filled).

A. GCN Average Metrics Across Folds for PLE and Transdiagnostic Samples

Connectome	Sample	Edge-Selection	Average Accuracy	Average Sensitivity	Average Specificity
Functional	PLE	SC-constrained	0.51 ± 0.06	0.53 ± 0.28	0.53 ± 0.28
Structural	PLE	Streamline Count ≥ 1	0.54 ± 0.06	0.63 ± 0.23	0.44 ± 0.24
Functional	Transdiagnostic	SC-constrained	0.55 ± 0.05	0.48 ± 0.23	0.63 ± 0.22
Structural	Transdiagnostic	Streamline Count ≥ 1	0.52 ± 0.07	0.49 ± 0.25	0.55 ± 0.19

B. Best GCN Model Metrics for PLE and Transdiagnostic Samples

Connectome	Sample	Edge-Selection	Accuracy	Sensitivity	Specificity
Functional	PLE	SC-constrained	0.62	0.38	0.87
Structural	PLE	Streamline Count ≥ 1	0.65	0.60	0.70
Functional	Transdiagnostic	SC-constrained	0.63	0.75	0.50
Structural	Transdiagnostic	Streamline Count ≥ 1	0.59	0.63	0.55

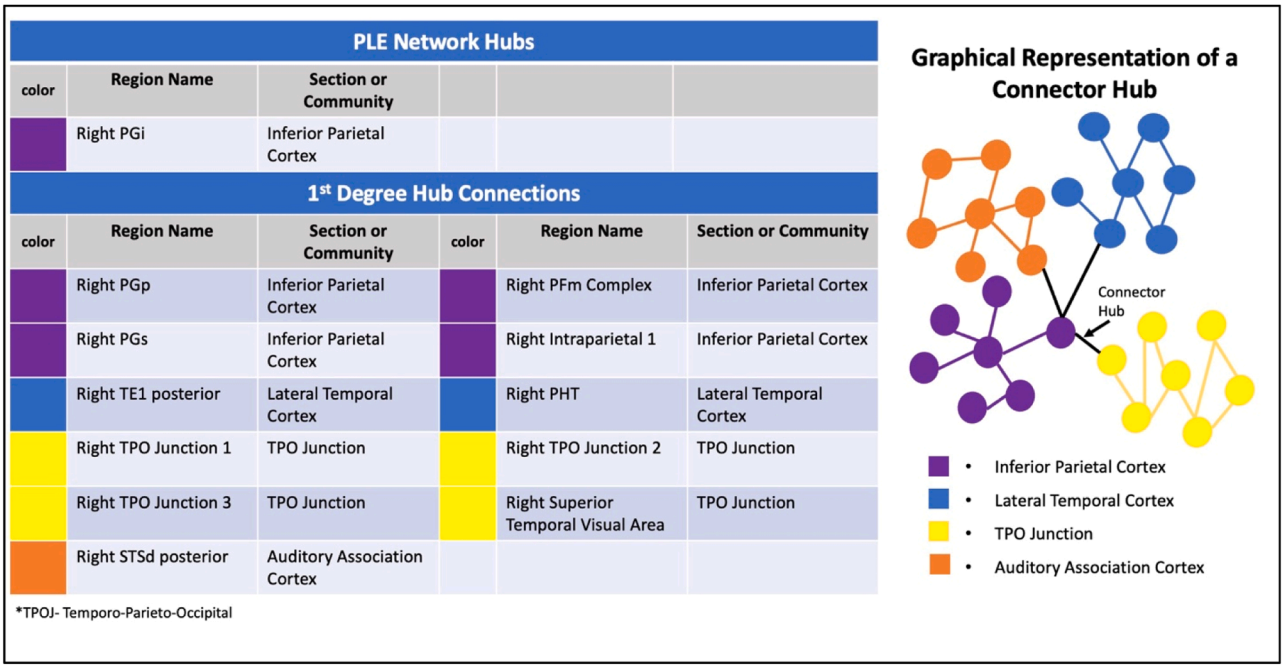

Fig. 3. List of hubs (top left) and their first-degree connections (bottom left). The hubs have connections with multiple other cortical sections making them connector hubs. A graphical representation of a connector hub is shown on the right with each color corresponding to a different cortical section defined by the Glasser parcellation.

similar to that in the transdiagnostic group which illustrates the predictive ability of SVMs and GCNs in a diagnostically heterogeneous transdiagnostic sample or symptomatically subsyndromal PLE sample. While this is a challenging classification problem, these findings are especially encouraging because our previous attempts to predict the group membership of familial high-risk individuals (who have up to 15% risk for conversion to psychoses) using traditional structural equation modeling with psychopathological, cognitive, and a global brain measure provided much lower accuracy and sensitivity (Eack et al., 2008; Shah et al., 2012). Although the classification metrics in our SVM and GCN models are not currently suitable for clinical deployment, it should be noted that these methods used only one type of predictor variable of SC or structurally constrained FC compared to our prior study that used multiple predictor variables in structural equation modeling. Further, even with these metrics, GCN models can highlight the importance of salient regions for classification or identify diagnostic groups in our best GCN model. Our study revealed a connector hub which contributes to the accuracy of classification. The model that provided the highest classification accuracy between PLE and controls used SC inputs. The PLE and control groups had a unique set of salient and common regions. The salient regions for PLE were in the dorsolateral prefrontal cortex, lateral temporal, and orbital and polar frontal cortex Glasser sections. Among the controls, the salient regions were in the auditory association, posterior cingulate, inferior frontal, paracentral lobular, mid cingulate, and superior parietal cortices Glasser sections. The salient regions common to both groups were in the primary visual, early visual, somatosensory, and motor, and subcortical sections. Adding other potential clinical variables or using more homogenous diagnostic groups may improve the classification accuracy, potentially making deep learning methods more appropriate for clinical applications.

Despite the relatively low accuracy of classification models observed in this study, our findings are valuable because of the possibility for early detection using machine learning that may help early interventions. Prior studies have demonstrated that early detection and interventions improve outcomes (Correll et al., 2018; de Koning et al.,

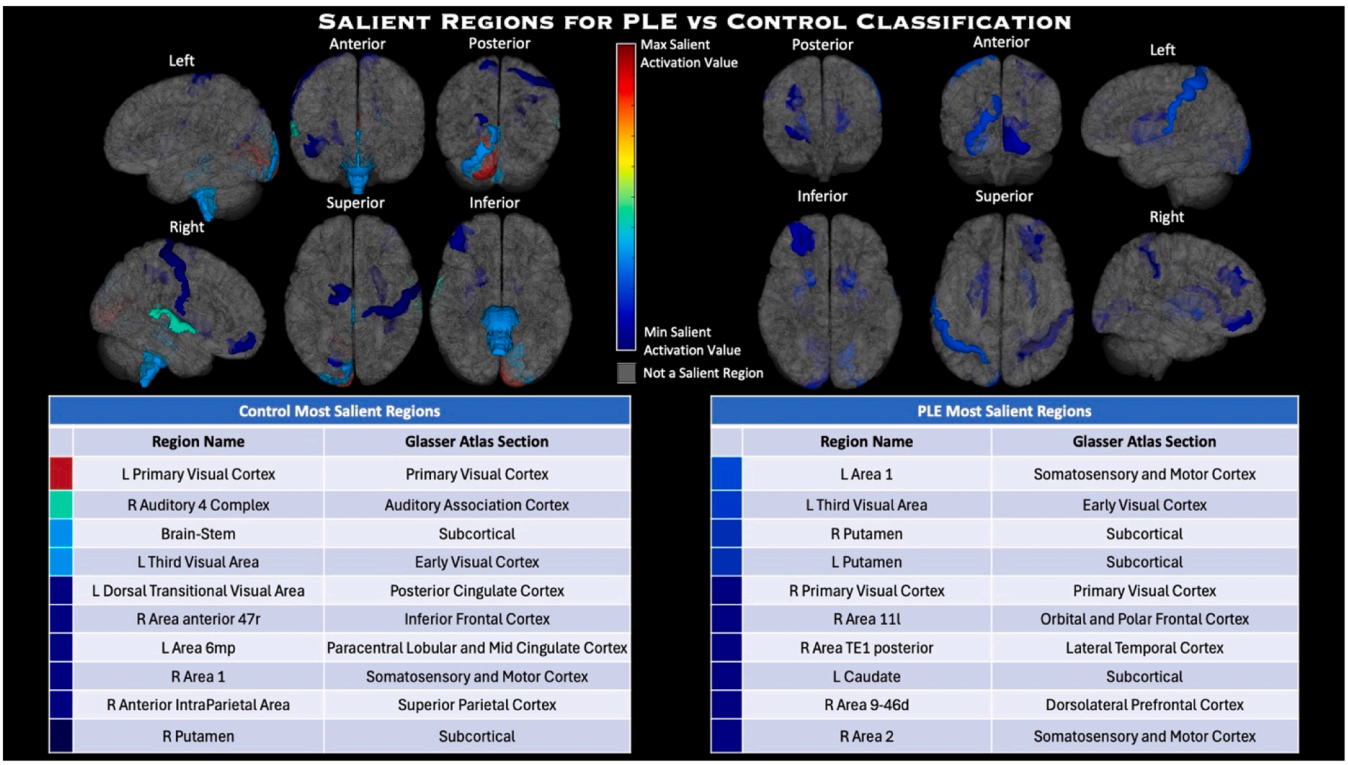

Control Most Salient Regions

Region Name	Glasser Atlas Section
L Primary Visual Cortex	Primary Visual Cortex
R Auditory 4 Complex	Auditory Association Cortex
Brain-Stem	Subcortical
L Third Visual Area	Early Visual Cortex
L Dorsal Transitional Visual Area	Posterior Cingulate Cortex
R Area anterior 47r	Inferior Frontal Cortex
L Area 6mp	Paracentral Lobular and Mid Cingulate Cortex
R Area 1	Somatosensory and Motor Cortex
R Anterior IntraParietal Area	Superior Parietal Cortex
R Putamen	Subcortical

PLE Most Salient Regions

Region Name	Glasser Atlas Section
L Area 1	Somatosensory and Motor Cortex
L Third Visual Area	Early Visual Cortex
R Putamen	Subcortical
L Putamen	Subcortical
R Primary Visual Cortex	Primary Visual Cortex
R Area 11l	Orbital and Polar Frontal Cortex
R Area TE1 posterior	Lateral Temporal Cortex
L Caudate	Subcortical
R Area 9-46d	Dorsolateral Prefrontal Cortex
R Area 2	Somatosensory and Motor Cortex

Fig. 4. Top 10 most salient regions for PLE and controls influencing classification. The control group (left) had the maximum activation value (red) and the minimum (the darkest blue). The PLE group (right) had a smaller range of activation values.

2009) and are cost-saving for people at high-risk of developing psychiatric disorders (Ologundudu et al., 2021). Further, we also found that certain interventions can not only benefit clinical outcomes but also improve neurobiological parameters (Eack et al., 2010; Keshavan et al., 2017). For these reasons, identification of people at high-risk for developing psychosis from psychiatrically healthy individuals is at the forefront of psychiatric research. Our study examines persons with a "soft" spectrum of psychopathology, namely people with PLE, who normally do not come under the care of a psychiatrist. Even if they do, no formal psychiatric diagnosis can be assigned leading to no clinically approved interventions. Attempts to classify disorders based on certain features of brain networks derived from MRI data is further useful because psychiatric disorders are defined entirely based on clusters of behaviors. The subjects in our dataset exist across a wide range of psychosis spectrum, and classifying across this spectrum is challenging because each disorder is shown to have different brain connectivity changes compared to a control subject.

Our analysis implemented two key machine learning tools, namely SVMs and GCNs. SVMs have a long history of application to classify groups and have well-understood theoretical foundations. However, SVMs use hyperplanes and learned parameters to classify brain networks that do not retain information about the network structure and can only handle one type of connectome input, either functional or structural (Miller et al., 2020). Since both structural and functional networks have been shown to be altered in people at high-risk for psychosis (Baker et al., 2014; Gong et al., 2017; Griffa et al., 2019; Prasad et al., 2023, 2015; Rikandi et al., 2022), traditional SVMs cannot handle differences in both structural and functional networks simultaneously. To address this limitation, we implemented the MultiVERSE embedding algorithm to combine SC and FC data before inputting into the SVM model. Further, we observed overfitting of SVM training models, possibly because of the high dimensionality of the networks. Even the MultiVERSE network embedding approach, which reduced the dimension in addition to combining the structural and functional data, did not fully address the problem of overfitting. Several issues remain regarding the use of MultiVERSE, such as deciding the number of iterations to run without a tolerance-based stopping criterion and the choice of embedding dimension. In separate simulation studies involving stochastic block model random graphs, the embedding dimension caused overfitting (Arroyo Relión et al., 2019), but in our data, it seems that the embedding dimension did not affect overfitting. Therefore, our future works will use regularized SVMs. In addition, since the SVMs do not retain the information about network structure, further analysis to identify network features that contributed to classification could not be implemented.

GCNs suited our connectome data due to their ability to retain nodal information of the input network structure and the use of three convolutional layers to model the graph inputs and distinguish the groups (Ktena et al., 2018; Lei et al., 2022). In this study, GCNs and SVMs provided similar classification results regardless of feature input. Some studies have reported a higher classification accuracy using GCNs on subjects with clinically diagnosed schizophrenia (Lei et al., 2022). Using our sample consisting of PLE with less prominent neurobiological changes compared to controls and a transdiagnostic sample with heterogenous networks, a classification accuracy of 63% is encouraging compared to structural equation modeling (SEM) in our previous studies (Eack et al., 2008; Shah et al., 2012). Our sample was balanced to avoid having the GCN models overfit the network features of one group compared to the other, unlike some other studies (Lei et al., 2022). An imbalanced sample can bias the model to provide higher balanced accuracy for the group with larger sample size since the model overfits that group because balanced accuracy is an average of sensitivity and specificity. Additionally, we were careful in entering the inputs in the proper order for the confusion matrix. We input the ground truth labels first and then the predicted labels next since entering in an incorrect order could yield a falsely high balanced accuracy by inflating sensitivity and/or specificity. Using a more homogeneous PLE sample, the classification accuracy improved which suggests that comparing two groups where

each group is relatively more homogeneous is better. This also shows that network connectivity of people at relatively higher risk for psychosis, namely PLE, and the transdiagnostic sample share similar features compared to controls.

The choice of classification method depends on the scientific question being asked and the resources available. While SVMs use relatively fewer computational resources and are faster to implement than GCNs, SVMs still provided the same accuracy when used on our samples (Supplemental Table 6). Previous studies have similarly reported high accuracy on other samples (Achalia et al., 2020; Gallo et al., 2023; Lei et al., 2022). However, SVMs may not be suitable for networks with a large number of nodes because of the risk of overfitting as we observed with the Glasser parcellation (377 cortical and subcortical regions) (Glasser et al., 2016). A study using the AAL atlas, which consisted of 90 nodes, showed better accuracy and did not mention overfitting in their data (Lei et al., 2022). Therefore, using a parcellation atlas with fewer regions may reduce overfitting in the SVMs but could lose valuable regional connectivity information provided by the Glasser atlas. In such instances, it may be preferable to use deep learning methods, such as GCNs implemented with careful consideration of the edge selection method, provided the computational resources are accessible. Our study shows that FC and SC inputs performed similarly most of the time, but the best classification accuracy was observed when the SC or structurally constrained FC was used. This observation indicates that integrating data from different modalities into a multiplex form could improve classification compared to inputting the monoplex networks. Further, GCNs can handle larger number of nodes supporting an additional advantage over SVMs.

In summary, we sought to classify PSD individuals consisting of PLE and a transdiagnostic sample with relatively high accuracy using SVMs and GCNs. We also identified that multiplexing SC and FC improves connectome classification while using structurally constrained FCs for GCNs resulting in marginal improvement. FC yielded 3–5% higher accuracy compared to other inputs on average for SVMs for the PSD and PLE samples, but for the transdiagnostic sample, MultiVERSE inputs were better with about 5% higher accuracy than other methods on average.

5. Conclusions

This paper used GCNs and SVMs to classify diverse psychosis samples of PSD, PLE, and transdiagnostic samples compared to healthy controls. Both methods produced similar accuracies of 65–69%. The edge selection method chosen for GCNs improved classification accuracy, suggesting that researchers should carefully choose biologically relevant edge selection methods. Our analysis also revealed that using a combination of FC and SC may improve classification, and we identified an inferior parietal region as hub in PLE subjects only. Because GCNs rely on network architecture, using GCNs permits further investigations using network analysis and CAM which, in our study, revealed hubs and salient regions unique to PLE subjects. In future studies to improve classification accuracy, we could use network features instead of the entire network as the input to reduce the number of features and implement feature selection methods that have been successful previously (Abdel Hady & Abd El-Hafeez, 2023, 2024). Future studies to improve classification accuracy could use clinical and cognitive scores as additional features within the machine learning models which have shown success in multi-disease classification using symptom descriptors (Hassan et al., 2024) and Hepatitis C prediction (Mamdouh Farghaly et al., 2023). Limitations of this study include a relatively small sample of the transdiagnostic group, but in future studies, we plan to replicate these results on publicly available subclinical or high-risk adolescents and young adults with high-risk of developing psychosis to further enhance the usefulness of the classification to develop early interventions and improve outcomes (Khosla, 2006; Kosztyán et al., 2022) and minimize long-term social and functional impairments. Additionally in future analysis, we will include soft classifiers which outputs group probabilities instead of a binary classification and are not present in this analysis. We will also include other dimensionality reduction techniques alongside SVMs such as principal component analysis, network-based nonparametric dimensionality reduction, and latent factor analysis (Khosla, 2006; Kosztyán et al., 2022; Nakayama et al., 2021).

Funding source

Funding was provided by NIMH grant numbers: R01MH115026 and R01MH112584 (KMP). Computational resources were provided by the Pittsburgh Supercomputing Center through ACCESS Discover grant BIO200047 (KMP). JC was supported in part by the National Science Foundation under grant NSF DMS 2413552.

Declaration of generative AI in scientific writing

No generative AI programs were used for this work, including the data analysis, drafting, and editing of the manuscript.

Code availability

The code for this analysis is publicly available on GitHub (https://github.com/madisonlewis323/PSD_GCN_analysis).

CRediT authorship contribution statement

Madison Lewis: Writing – review & editing, Writing – original draft, Validation, Formal analysis, Data curation. **Wenlong Jiang:** Writing – review & editing, Writing – original draft, Formal analysis. **Nicholas D. Theis:** Writing – review & editing, Writing – original draft, Resources, Data curation. **Joshua Cape:** Writing – review & editing, Writing – original draft, Validation, Supervision. **Konasale M. Prasad:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Project administration, Methodology, Investigation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data is available through the UK Biobank

Acknowledgements

We thank the UK Biobank for providing access to their database for all the multimodal data used in the analysis. In addition, we would like to extend thanks to Du Lei and other authors for helping in the Class Activation Mapping implementation.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.neunet.2024.106771.

References

Abdel Hady, D. A., & Abd El-Hafeez, T. (2023). Predicting female pelvic tilt and lumbar angle using machine learning in case of urinary incontinence and sexual dysfunction. *Scientific Reports*, 13.

Abdel Hady, D. A., & Abd El-Hafeez, T. (2024). Revolutionizing core muscle analysis in female sexual dysfunction based on machine learning. *Scientific Reports*, 14.

Achalia, R., Sinha, A., Jacob, A., Achalia, G., Kaginalkar, V., Venkatasubramanian, G., et al. (2020). A proof of concept machine learning analysis using multimodal neuroimaging and neurocognitive measures as predictive biomarker in bipolar disorder. *Asian Journal of Psychiatry, 50*, Article 101984.

Arroyo Relión, J. D., Kessler, D., Levina, E., & Taylor, S. F. (2019). Network classification with applications to brain connectomics. *The Annals of Applied Statistics, 13*, 1648–1677.

Arslan, S., Ktena, S.I., Glocker, B., & Rueckert, D. (2018). Graph saliency maps through spectral convolution networks: Application to sex classification with brain connectivity. In: ArXiv.

Baker, J. T., Holmes, A. J., Masters, G. A., Yeo, B. T., Krienen, F., Buckner, R. L., et al. (2014). Disruption of cortical association networks in schizophrenia and psychotic bipolar disorder. *JAMA Psychiatry, 71*, 109–118.

Bassett, D. S., Bullmore, E., Verchinski, B. A., Mattay, V. S., Weinberger, D. R., & Meyer-Lindenberg, A. (2008). Hierarchical organization of human cortical networks in health and schizophrenia. *The Journal of Neuroscience, 28*, 9239–9248.

Beckmann, C. F., & Smith, S. M. (2004). probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE Transactions on Medical Imaging, 23*, 137–152.

Bordier, C., Nicolini, C., & Bifone, A. (2017). Graph analysis and modularity of brain functional connectivity networks: Searching for the optimal threshold. *Frontiers in Neuroscience, 11*, 441.

Buitrago, P., & Nystrom, N. (2021). *Neocortex and bridges-2: A high performance AI+HPC ecosystem for science, discovery, and societal good* (pp. 205–219). High Performance Computing.

Chu, Y. J., & Liu, T. H. (1965). On the shortest arborescence of a directed graph. *Science Sinica, 14*, 1396–1400.

Cohen, J. R., & D'Esposito, M. (2016). The segregation and integration of distinct brain networks and their relationship to cognition. *The Journal of Neuroscience, 36*, 12083–12094.

Correll, C. U., Galling, B., Pawar, A., Krivko, A., Bonetto, C., Ruggeri, M., et al. (2018). Comparison of early intervention services vs treatment as usual for early-phase psychosis: A systematic review, meta-analysis, and meta-regression. *JAMA Psychiatry, 75*, 555–565.

Cover, T., & Hart, P. (1967). *Nearest neighbor pattern classification., 13* pp. 21–27). IEEE Transactions on Information Theory.

de Koning, M. B., Bloemen, O. J., van Amelsvoort, T. A., Becker, H. E., Nieman, D. H., van der Gaag, M., et al. (2009). Early intervention in patients at ultra high risk of psychosis: Benefits and risks. *Acta Psychiatrica Scandinavica, 119*, 426–442.

Eack, S. M., Hogarty, G. E., Cho, R. Y., Prasad, K. M., Greenwald, D. P., Hogarty, S. S., et al. (2010). Neuroprotective effects of cognitive enhancement therapy against gray matter loss in early schizophrenia: Results from a 2-year randomized controlled trial. *Arch Gen Psychiatry, 67*, 674–682.

Eack, S. M., Prasad, K. M., Montrose, D. M., Goradia, D. D., Dworakowski, D., Miewald, J., et al. (2008). An integrated psychobiological predictive model of emergent psychopathology among young relatives at risk for schizophrenia. *Progress in Neuro-Psychopharmacology & Biological Psychiatry, 32*, 1873–1878.

Edmonds, J. (1967). Optimum branchings. Journal of RESEARCH of the National Bureau of Standards, 71B, 233–240.

Farahani, F. V., Karwowski, W., & Lighthall, N. R. (2019). Application of graph theory for identifying connectivity patterns in human brain networks: A systematic review. *Frontiers in Neuroscience, 13*, 585.

Fischl, B. (2012). FreeSurfer. *NeuroImage, 62*, 774–781.

Gallo, S., El-Gazzar, A., Zhutovsky, P., Thomas, R. M., Javaheripour, N., Li, M., et al. (2023). Functional connectivity signatures of major depressive disorder: Machine learning analysis of two multicenter neuroimaging studies. *Molecular Psychiatry, 28*, 3013–3022.

Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., et al. (2016). A multi-modal parcellation of human cerebral cortex. *Nature, 536*, 171–178.

Gong, Q., Hu, X., Pettersson-Yeo, W., Xu, X., Lui, S., Crossley, N., et al. (2017). Network-level dysconnectivity in drug-naïve first-episode psychosis: Dissociating transdiagnostic and diagnosis-specific alterations. *Neuropsychopharmacology : Official Publication of the American College of Neuropsychopharmacology, 42*, 933–940.

Griffa, A., Baumann, P. S., Klauser, P., Mullier, E., Cleusix, M., Jenni, R., et al. (2019). Brain connectivity alterations in early psychosis: From clinical to neuroimaging staging. *Translational Psychiatry, 9*, 62.

Griffanti, L., Salimi-Khorshidi, G., Beckmann, C. F., Auerbach, E. J., Douaud, G., Sexton, C. E., et al. (2014). ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging. *NeuroImage, 95*, 232–247.

Gutmann, M., & Hyvärinen, A. (2010). Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *, 9. Proceedings of the 13th international conference on artificial intelligence and statistics*.

Hassan, E., Abd El-Hafeez, T., & Shams, M. Y (2024). Optimizing classification of diseases through language model analysis of symptoms. *Scientific Reports*, 14.

Jaaskelainen, E., Juola, P., Hirvonen, N., McGrath, J. J., Saha, S., Isohanni, M., et al. (2013). A systematic review and meta-analysis of recovery in schizophrenia. *Schizophrenia Bulletin, 39*, 1296–1306.

Jenkinson, M., Beckmann, C. F., Behrens, T. E., Woolrich, M. W., & Smith, S. M. (2012). Fsl. *NeuroImage, 62*, 782–790.

Kang, S. (2021). *K-nearest neighbor learning with graph neural networks, 9*. Mathematics.

Kaymaz, N., Drukker, M., Lieb, R., Wittchen, H. U., Werbeloff, N., Weiser, M., et al. (2012). Do subthreshold psychotic experiences predict clinical outcomes in unselected non-help-seeking population-based samples? A systematic review and meta-analysis, enriched with new results. *Psychological Medicine, 42*, 2239–2253.

Keshavan, M. S., Eack, S. M., Prasad, K. M., Haller, C. S., & Cho, R. Y. (2017). Longitudinal functional brain imaging study in early course schizophrenia before and after cognitive enhancement therapy. *NeuroImage, 151*, 55–64.

Khlifi, M. K., Boulila, W., & Farah, I. R. (2023). Graph-based deep learning techniques for remote sensing applications: Techniques, taxonomy, and applications — A comprehensive review. *Computer Science Review*, 50.

Khosla, N. (2006). *Dimensionality reduction using factor analysis*. Queensland, Australia: Griffith University.

Kosztyán, Z. T., Kurbucz, M. T., & Katona, A. I. (2022). *Network-based dimensionality reduction of high-dimensional, low-sample-size datasets, 251*. Knowledge-Based Systems.

Ktena, S. I., Parisot, S., Ferrante, E., Rajchl, M., Lee, M., Glocker, B., et al. (2018). Metric learning with spectral graph convolutions on brain connectivity networks. *NeuroImage, 169*, 431–442.

Kurtzer, G., Sochat, V., & Bauer, M. (2017). Singularity: Scientific containers for mobility of compute. *PloS one*.

Laursen, T. M., Munk-Olsen, T., & Vestergaard, M. (2012). Life expectancy and cardiovascular mortality in persons with schizophrenia. *Current Opinion in Psychiatry, 25*, 83–88.

Lei, D., Qin, K., Pinaya, W. H. L., Young, J., van Amelsvoort, T., Marcelis, M., et al. (2022). Graph convolutional networks reveal network-level functional dysconnectivity in schizophrenia. *Schizophrenia Bulletin*.

Lewis, M., Santini, T., Theis, N., Muldoon, B., Dash, K., Rubin, J., et al. (2023). Modular architecture and resilience of structural covariance networks in first-episode antipsychotic-naive psychoses. *Scientific Reports, 13*, 7751.

Li, G. (2024). *Maximum weight spanning tree (Undirected)*. MATLAB Central File Exchange.

Linscott, R. J., & van Os, J. (2013). An updated and conservative systematic review and meta-analysis of epidemiological evidence on psychotic experiences in children and adults: On the pathway from proneness to persistence to dimensional expression across mental disorders. *Psychological Medicine, 43*, 1133–1149.

Littlejohns, T. J., Holliday, J., Gibson, L. M., Garratt, S., Oesingmann, N., Alfaro-Almagro, F., et al. (2020). The UK Biobank imaging enhancement of 100,000 participants: Rationale, data collection, management and future directions. *Nature Communications, 11*, 2624.

Mamdouh Farghaly, H., Shams, M. Y., & Abd El-Hafeez, T. (2023). Hepatitis C Virus prediction based on machine learning framework: A real-world case study in Egypt. *Knowledge of Information Systems, 65*, 2595–2617.

McIntosh, A. M., Whalley, H. C., McKirdy, J., Hall, J., Sussmann, J. E., Shankar, P., et al. (2008). Prefrontal function and activation in bipolar disorder and schizophrenia. *The American Journal of Psychiatry, 165*, 378–384.

Miller, C. H., Sacchet, M. D., & Gotlib, I. H. (2020). *Support vector machines and affective science, 12* pp. 297–308). Emotion Review.

Miller, K. L., Alfaro-Almagro, F., Bangerter, N. K., Thomas, D. L., Yacoub, E., Xu, J., et al. (2016). Multimodal population brain imaging in the UK Biobank prospective epidemiological study. *Nature Neuroscience, 19*, 1523–1536.

Mills, K. (2016). HCP-MMP1.0 projected on fsaverage. In figshare (Ed.).

Nakayama, Y., Yata, K., & Aoshima, M. (2021). Clustering by principal component analysis with Guassian kernel in high-dimension, low-sample-size settings. *Journal of Multivariate Statistics, 185*.

Neurolab, C. (2017). HCP-MMP1.0 volumetric (NIfTI) masks in native structural space. In (pp. Dataset). figshare.

Ologundudu, O. M., Lau, T., Palaniyappan, L., Ali, S., & Anderson, K. K. (2021). Interventions for people at ultra-high risk for psychosis: A systematic review of economic evaluations. *Early Intervention in Psychiatry, 15*, 1115–1126.

Pereira, F., Mitchell, T., & Botvinick, M. (2009). Machine learning classifiers and fMRI: A tutorial overview. *NeuroImage, 45*, S199–S209.

Pio-Lopez, L., Valdeolivas, A., Tichit, L., Remy, É., & Baudot, A. (2021). MultiVERSE: A multiplex and multiplex-heterogeneous network embedding approach. *Scientific Reports, 11*, 8794.

Prasad, K. M. (2017). Course, prognosis, and outcome of schizophrenia and related disorders. M. J. Marcsisin, J. M. Gannon & J. B. Rosenstock (Eds.). *Schizophrenia and related disorders* (pp. 111–142). New York: Oxford University Press.

Prasad, K. M., Muldoon, B., Theis, N., Iyengar, S., & Keshavan, M. S. (2023). Multipronged investigation of morphometry and connectivity of hippocampal network in relation to risk for psychosis using ultrahigh field MRI. *Schizophrenia Research, 256*, 88–97.

Prasad, K. M., Upton, C. H., Schirda, C. S., Nimgaonkar, V. L., & Keshavan, M. S. (2015). White matter diffusivity and microarchitecture among schizophrenia subjects and first-degree relatives. *Schizophrenia Research, 161*, 70–75.

Rikandi, E., Mäntylä, T., Lindgren, M., Kieseppä, T., Suvisaari, J., & Raij, T. T. (2022). Functional network connectivity and topology during naturalistic stimulus is altered in first-episode psychosis. *Schizophrenia Research, 241*, 83–91.

Robinson, D. G., Woerner, M. G., McMeniman, M., Mendelowitz, A., & Bilder, R. M. (2004). Symptomatic and functional recovery from a first episode of schizophrenia or schizoaffective disorder. *The American Journal of Psychiatry, 161*, 473–479.

Rubinov, M., & Sporns, O. (2010). Complex network measures of brain connectivity: Uses and interpretations. *NeuroImage, 52*, 1059–1069.

Salimi-Khorshidi, G., Douaud, G., Beckmann, C. F., Glasser, M. F., Griffanti, L., & Smith, S. M. (2014). Automatic denoising of functional MRI data: Combining independent component analysis and hierarchical fusion of classifiers. *NeuroImage, 90*, 449–468.

Shah, J., Eack, S. M., Montrose, D. M., Tandon, N., Miewald, J. M., Prasad, K. M., et al. (2012). Multivariate prediction of emerging psychosis in adolescents at high risk for schizophrenia. *Schizophrenia Research, 141*, 189–196.

Smith, S.M., Alfaro-Almagro, F., & Miller, K.L. (2022). UK Biobank Brain Imaging Documentation. In.

Smucny, J., Tully, L. M., Howell, A. M., Lesh, T. A., Johnson, S. L., O'Reilly, R. C., et al. (2021). Schizophrenia and bipolar disorder are associated with opposite brain reward anticipation-associated response. *Neuropsychopharmacology : Official Publication of the American College of Neuropsychopharmacology, 46*, 1152–1160.

Theis, N., Bahuguna, J., Rubin, J.E., Cape, J., Iyengar, S., & Prasad, K.M. (2024). Subject-specific maximum entropy model of resting state fMRI shows diagnostically distinct patterns of energy state distributions. bioRxiv.

Theis, N., Rubin, J., Cape, J., Iyengar, S., & Prasad, K. M. (2023). Threshold selection for brain connectomes. *Brain Connectivity, 13*, 383–393.

Tournier, J. D., Smith, R., Raffelt, D., Tabbara, R., Dhollander, T., Pietsch, M., et al. (2019). MRtrix3: A fast, flexible and open software framework for medical image processing and visualisation. *NeuroImage, 202*, Article 116137.

Welham, J., Scott, J., Williams, G., Najman, J., Bor, W., O'Callaghan, M., et al. (2009). Emotional and behavioural antecedents of young adults who screen positive for non-affective psychosis: A 21-year birth cohort study. *Psychological Medicine, 39*, 625–634.

Wigman, J. T., van Nierop, M., Vollebergh, W. A., Lieb, R., Beesdo-Baum, K., Wittchen, H. U., et al. (2012). Evidence that psychotic symptoms are prevalent in disorders of anxiety and depression, impacting on illness onset, risk, and severity–implications for diagnosis and ultra-high risk research. *Schizophrenia Bulletin, 38*, 247–257.

Wikström, A., Tuulio-Henriksson, A., Perälä, J., Saarni, S., & Suvisaari, J. (2015). Psychotic like experiences (PLE's) in middle-aged adults. *Schizophrenia Research, 169*, 313–317.

Yang, S., Zhu, F., Ling, X., Liu, Q., & Zhao, P. (2021). Intelligent health care: Applications of deep learning in computational medicine. *Frontiers in Genetics, 12*, Article 607471.

Yu, M., Linn, K. A., Cook, P. A., Phillips, M. L., McInnis, M., Fava, M., et al. (2018). Statistical harmonization corrects site effects in functional connectivity measurements from multi-site fMRI data. *Human Brain Mapping, 39*, 4213–4227.

Ziermans, T. B. (2013). Working memory capacity and psychotic-like experiences in a general population sample of adolescents and young adults. *Frontiers in Psychiatry, 4*, 161.