Unveiling Brain-Body Axis Interactions in Psychiatric-Endocrine Disorder with Deep Graph Causal Neural Network

Zhonghua Wan[1†], Chengzhe Zhang[1†], Yu Xie[1], Yifei He[1], Qingwei You[1], Peng Wang[1], Jiaolong Qin[1], Huifeng Zhang[2], and Ye Wu[1✉]

[1] School of Computer Science and Engineering, Nanjing University of Science and Technology, Nanjing, China
wuye@njust.edu.cn

[2] Division of Mood Disorders, Shanghai Mental Health Center, Shanghai Jiao Tong University School of Medicine, Shanghai, China

Abstract. Understanding the role of the brain-body axis in the causal relationship between psychiatric disorders and endocrine disorders is crucial for advancing multi-organ disease modeling and precision medicine. Despite the importance of understanding these interactions, the development of a mature framework has been hindered by traditional models' inability to move beyond mere statistical associations to uncover true causal directionality in high-dimensional, population-scale data. In this study, we proposed a graph neural network (GNN)-based framework to construct causal graphs capturing interactions among brain imaging features, organ-level physiological indicators, and disease outcomes. Leveraging the directed acyclic graph (DAG)-GNN model, we perform causal structure learning on the comprehensive UK Biobank cohort, enabling the discovery of DAGs that reflect potential causal dependencies. Experimental results demonstrate that our method effectively uncovers interpretable brain-body-disease pathways, revealing shared white-matter vulnerabilities across distinct disorders and highlighting how intermediate physiological systems mediate these interactions, providing insights into psychiatric and endocrine conditions. Our code is available at https://github.com/NIMGroup/UKBCausal.

Keywords: Brain-body axis · Causal interaction · Psychiatric-endocrine disorder

1 Introduction

The human body operates as a highly interconnected system, in which the brain and body organs maintain a continuous physiological dialogue through complex pathways, often referred to as the brain-body axis [13]. Emerging evidence

† These authors contributed equally to this work.
✉ Corresponding author: Ye Wu.

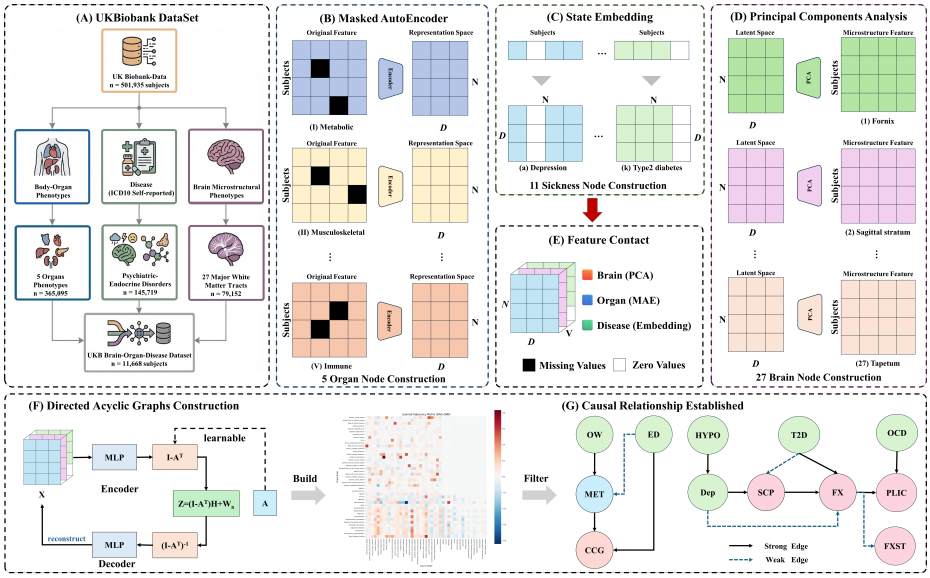

Fig. 1. The overview of our proposed framework.

suggests that this interaction plays a vital role in the progression of various systemic diseases. For instance, metabolic dysfunction can significantly accelerate neurodegeneration, and psychiatric and endocrine disorders show a bidirectional association whose causal directionality and brain–body pathways remain unclear [3]. Consequently, analyzing the intricate relationships between brain imaging phenotypes (e.g., white matter integrity) and organ-level physiological indicators is essential for achieving a holistic understanding of disease mechanisms and advancing precision medicine.

While robust statistical associations within the brain-body network are well documented, deciphering their underlying causal mechanisms is severely hindered by the inability of traditional models to process high-dimensional, incomplete, multimodal data [5, 21]. Previous studies have consistently reported shared white-matter alterations across major psychiatric disorders [10, 11, 22] and have highlighted bidirectional links between metabolic/endocrine dysfunction and mental health [3, 24]. More recently, comprehensive population-scale analyses have further underscored the systemic nature of these interactions, revealing how multi-organ health and multimorbidity risks are deeply intertwined [20, 19, 7, 4]. However, these foundational findings primarily evaluate relationships in isolation and rely on correlations, which fail to reveal the directionality of influence captured by Directed Acyclic Graphs (DAGs) [23]. Distinguishing whether brain structural changes precede organ dysfunction or vice versa is critical for identifying early biomarkers and therapeutic targets. Furthermore, real-world population datasets are invariably incomplete, leading to missing values that can severely bias structure-learning algorithms [12]. Therefore, there

is an urgent need for a unified framework that can handle missing data while effectively recovering interpretable causal graphs from high-dimensional observational biomedical data [6].

In this study, we present a novel framework for understanding the causal relationships among the brain, body organs, and diseases. Unlike previous approaches that primarily evaluate brain-body interactions in isolated pairs and rely strictly on statistical correlations, our method establishes a unified causal network that simultaneously captures the directionality of influences among brain microstructural features, multi-organ physiological states, and diverse disease outcomes. Our work comprises three main contributions: First, we used dMRI values [2] of major brain fiber bundles and feature values of major body organs from the UKB dataset to construct DAGs for exploration. Second, to address missing data in the UKB dataset, we employed a masked autoencoder [8, 14], a method distinct from traditional mean imputation, to impute missing values. Finally, we applied the DAG-GNN framework to construct an interpretable causal graph from observational data, revealing the underlying directional dynamics among 27 regional white matter features, 5 key physiological systems, and 11 distinct psychiatric and endocrine disorders.

2 Method

2.1 Dataset

We utilized data from the UK Biobank [18] to identify participants with psychiatric disorders by integrating hospital inpatient records and self-reported medical histories. We focused on seven major psychiatric disorders: schizophrenia, depression, bipolar disorder, anxiety disorders, eating disorders, obsessive-compulsive disorder, and adjustment disorders, as well as four endocrine disorders: hypothyroidism, hyperthyroidism, type 2 diabetes, and being overweight. To assess the physiological states of body organs, we followed [20] and extracted relevant phenotypes across the musculoskeletal, immune, metabolic, renal, and hepatic systems. Musculoskeletal health was evaluated with anthropometric measures and heel bone mineral density, while immune function was assessed through blood cell counts and C-reactive protein levels. Metabolic profiles included lipid levels and HbA1c; renal and hepatic function were assessed using creatinine, albumin, ALT, and AST. We also captured brain microstructural integrity using diffusion MRI data processed via the UKB-brain-pipeline [2], extracting nine tensor-derived and NODDI metrics averaged across 27 major white matter tracts from the JHU ICBM-DTI-81 atlas [15], resulting in a high-dimensional feature vector of the brain's communication network.

2.2 Masked Autoencoder for Organ-Level Representation Learning

In large-scale cohorts such as the UK Biobank, organ-level physiological phenotypes are often incomplete due to varying examination protocols and subject compliance [12]. Consequently, each subject may have only a partial set of

features, leading to substantial missingness. To ensure reliable data while maintaining sample size, we used an organ-specific subject-filtering strategy, retaining subjects with at least 70% of observed features per organ. The final study cohort was formed by intersecting subjects across all organ systems, creating a consistent set for downstream analysis.

Let $X^{(o)} \in \mathbb{R}^{N \times D_o}$ denote the feature matrix of organ O, where N is the number of subjects and D_o is the number of organ-specific phenotypes rather than explicitly imputing missing entries in X_o, which may introduce strong modeling assumptions and propagate bias into downstream analyses, we aim to learn a compact latent representation that is robust to partial observations. We adopt a masked autoencoder (MAE) [8] framework tailored to organ-level data. For each organ O, we restrict training to subjects with complete observations and synthetically introduce missingness by randomly masking a subset of features. Formally, for each subject i, we sample a binary mask vector $m_i \in \{0,1\}^{D_o}$, where M_{ij} indicates that the j-th feature is masked. The masked input is then given by $\tilde{x}_i^{(o)} = m_i \odot x_i^{(o)}$, where $\odot$ denotes element-wise multiplication. The masked autoencoder consists of an encoder $f_\theta^{(o)}$ that maps the partially observed input to a latent representation $z_i^{(o)} = f_\theta^{(o)}\left(\tilde{x}_i^{(o)}\right)$, followed by a decoder $g_\theta^{(o)}$ that reconstructs the original feature vector $\hat{x}_i^{(o)} = g_\theta^{(o)} \odot z_i^{(o)}$.

The training objective is defined exclusively on the masked entries, encouraging the model to infer missing values from observed context while avoiding trivial identity mappings:

$$L_{\text{MAE}}^{(o)} = \frac{1}{N} \sum_{i=1}^{N} \left\| (1 - m_i) \odot \left(x_i^{(o)} - \hat{x}_i^{(o)} \right) \right\|_2^2, \tag{1}$$

The masked autoencoder captures the low-dimensional structure of organ-level physiological profiles by reconstructing masked features from incomplete inputs, without requiring imputation of unobserved values. The learned latent representation $z_i^{(o)}$ acts as a robust organ-specific embedding for our causal graph discovery framework.

2.3 Unified Node Representation for Brain Graph Construction

To create a unified causal graph that connects brain regions, body organs, and diseases, all nodes must share a consistent feature dimensionality. Unlike organ-level phenotypes, brain dMRI features from the UK Biobank are fully observed, as all subjects underwent dMRI at imaging instance 2. Standardized preprocessing pipelines were used to extract microstructural metrics, including fractional anisotropy, anisotropy mode, and principal diffusivities.

For each tract r, we denote the corresponding dMRI feature matrix as $X^{(r)} \in \mathbb{R}^{N \times D}$, where N is the number of subjects and D_r is the number of region-specific diffusion metrics. Although these features are complete, their dimensionality varies across regions and is generally higher than that of organ-level

latent representations. Directly incorporating these heterogeneous feature spaces into a unified causal model would introduce scale imbalance and hinder graph learning. To address this issue, we apply principal component analysis (PCA) independently for each brain region [1] to obtain compact, comparable node-level representations. Specifically, for each region r, PCA is performed on the standardized feature matrix $X(r)$, yielding a low-dimensional embedding

$$Z^{(r)} = X^{(r)}W^{(r)}, \quad W^{(r)} \in \mathbb{R}^{D_r \times d}, \tag{2}$$

where d denotes the target embedding dimension. The value of d is chosen to match the latent dimensionality of organ-level nodes learned by the masked autoencoder, thereby ensuring representational consistency across node types.

The resulting embedding $z_i^{(r)} \in \mathbb{R}^d$ serves as the brain-region node feature for subject i. By aligning the dimensionality of brain and organ representations, we enable their seamless integration into a unified node-feature tensor that serves as input to subsequent directed acyclic graph construction.

2.4 Directed Acyclic Graphs Construction

Using unified node representations, we infer DAGs that capture causal dependencies among brain regions, body organs, and disease phenotypes. Let V represent the total number of nodes, including brain, organ, and disease nodes. For each subject i, we create a node feature matrix $X_i \in \mathbb{R}^{V \times d}$, with d being the shared latent dimensionality derived from PCA, masked autoencoders, and disease-specific representations. Our goal is to learn a weighted adjacency matrix $A \in \mathbb{R}^{V \times V}$, where A_{uv} represents the directed influence from node u to node v, while maintaining acyclicity for causal interpretability.

We employ a variational DAG-GNN framework [23] to learn node interactions and graph structure. A graph encoder parameterized by ϕ maps node features to a latent edge distribution $q_\phi(Z, A \mid X)$, where Z captures node-wise dependencies and A is the learned adjacency matrix. A decoder reconstructs the observed node features from Z and A. The model is trained by maximizing the evidence lower bound (ELBO):

$$L_{\text{ELBO}} = \mathbb{E}_{q_\phi(Z,A|X)}[\log p_\theta(X \mid Z, A)] - \text{KL}(q_\phi(Z, A \mid X) \,\|\, p(Z, A)), \tag{3}$$

where p denotes the decoder likelihood and $p(Z, A)$ is a sparsity-inducing prior over graph structures. To explicitly enforce the DAG property, we incorporate a differentiable acyclicity constraint on the adjacency matrix following prior work [25]. Specifically, the constraint is defined as

$$h(A) = \text{tr}\big(\exp(A \odot A)\big) - V = 0, \tag{4}$$

where $\odot$ denotes element-wise multiplication. This formulation ensures that $h(A) = 0$ if and only if A corresponds to a directed acyclic graph. The overall training objective combines the ELBO with the acyclicity constraint via an augmented Lagrangian:

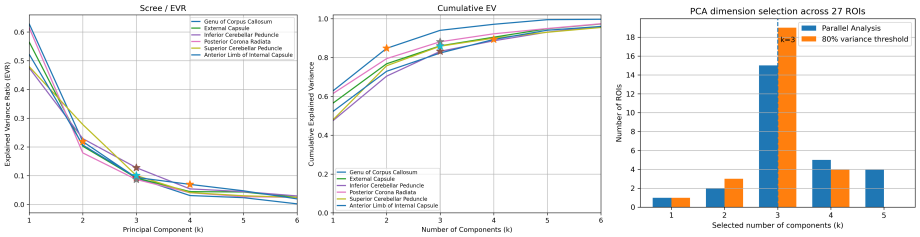

Fig. 2. PCA-Based Brain ROI Dimension Selection

$$L = L_{\text{ELBO}} + \lambda h(A) + \frac{\rho}{2} h(A)^2 + \tau \|A\|_1, \tag{5}$$

where λ and ρ are Lagrange multipliers, and the L penalty promotes sparsity in the learned graph.

After training, the learned adjacency matrix A is thresholded to create a sparse causal graph. Edges with small absolute weights are removed, whereas those with larger weights indicate stronger dependencies. This graph is acyclic by construction and represents the estimated causal relationships among brain regions, organs, and diseases.

3 Experiments and Results

3.1 PCA-Based Selection of Brain ROI Embedding Dimension

To select a unified embedding dimension for brain ROIs in our graph model, we performed PCA-based dimensionality selection [9] on the dMRI features of 27 white-matter tract ROIs. Features were standardized and decomposed into principal components. We analyzed the explained variance ratio (EVR) and cumulative explained variance (Cumulative EV) as functions of the number of components k. As shown in Fig. 2 (left), variance is mostly concentrated in the first few components, with EVR dropping sharply after the first component and the cumulative curve approaching a plateau. We used two criteria for selection: (i) an 80% cumulative variance rule and (ii) Parallel Analysis (PA) to retain components with eigenvalues above the 95th percentile from Gaussian reference data (star markers in Fig. 2). The results (Fig. 2, right) indicate a consensus on $k = 3$, suggesting that three components capture most of the variance across most ROIs. Thus, we set the dimension of the brain ROI embedding to $d = 3$ for consistency across experiments.

3.2 MAE can effectively characterize missing features

We assessed robustness to missing features through an extra-masking reconstruction evaluation on organ-specific feature matrices. Subjects were randomly split

Table 1. Extra-masking reconstruction RMSE on a 20% test split with 30% masking.

Organ	Masked AutoEncoder	Mean+PCA	SoftImpute
Hepatic	**0.772 ± 0.008**	0.827 ± 0.006	0.828 ± 0.007
Immune	0.731 ± 0.004	0.894 ± 0.004	**0.594 ± 0.005**
Metabolic	**0.642 ± 0.003**	0.771 ± 0.002	0.706 ± 0.003
Musculoskeletal	**0.743 ± 0.003**	0.854 ± 0.003	0.780 ± 0.003
Renal	**0.799 ± 0.004**	0.880 ± 0.006	0.873 ± 0.006

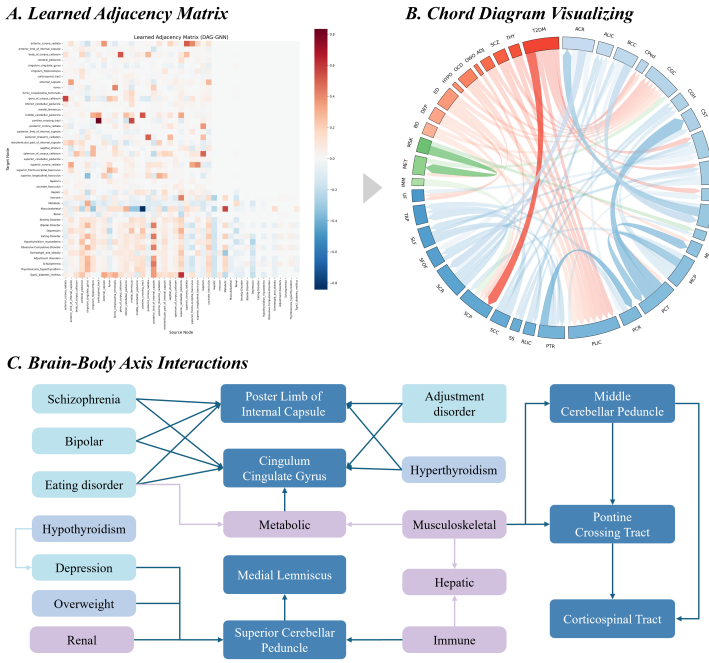

Fig. 3. Causal Graph Learned by DAG-GNN.

into training and test sets with an 80-to-20 ratio. Features were standardized using a StandardScaler from the training set and applied to both sets. We simulated missingness by masking 30% of the test-set entries and measured performance as the RMSE over the masked positions. The procedure was repeated 5 times, and the results were reported as mean ± standard deviation. Results in Table 1 demonstrate that MAE consistently outperforms the Mean+PCA baseline across all organ systems. MAE achieves the lowest reconstruction error across most categories compared with SoftImpute. By leveraging deep non-linear representations rather than global low-rank assumptions, MAE effectively recovers missing values in high-dimensional physiological phenotypes. While SoftImpute performs slightly better on immune features, MAE demonstrates the best overall robustness, justifying its adoption for reliable graph node construction.

3.3 Causal discovery of brain–organ–disease links

Leveraging the completed multi-node features, we applied the DAG-GNN architecture to infer the directed acyclic graph (DAG) underlying brain–organ–disease interactions. By thresholding the learned weighted adjacency matrix to retain the most essential edges ($w \succ 0.5$) and the second most important edges ($w \succ 0.2$), the causal structure visualized in Fig. 3 is obtained.

Disease-rooted subgraph analysis revealed distinct topological patterns in the relationship between systemic conditions and brain white-matter microstructure. First, we observed convergent vulnerability across two major tracts—the cingulum within the cingulate gyrus and the posterior limb of the internal capsule —both of which repeatedly emerged as downstream targets. These pathways received direct causal edges from multiple psychiatric disorders (e.g., schizophrenia, bipolar disorder, obsessive–compulsive disorder, eating disorder, and adjustment disorder) as well as endocrine dysfunction (hyperthyroidism), suggesting shared neuroanatomical endpoints across heterogeneous disease domains [10, 11]. This convergence highlights their vulnerability as critical network hubs and aligns with established transdiagnostic structural signatures [17]. Second, the inferred graph highlighted mediated pathways that distinguish direct tract-level effects from systemic cascades. Being overweight indirectly affected the cingulum via metabolic dysfunction (overweight $\rightarrow$ metabolic dysfunction $\rightarrow$ cingulum), and hypothyroidism influenced the superior cerebellar peduncle (SCP) via depression (hypothyroidism $\rightarrow$ depression $\rightarrow$ SCP). This systematic cascading effect is consistent with the existing literature [16], which indicates that obesity and metabolic syndrome primarily drive structural changes in the white matter of the limbic system, including the cingulate. Besides disease-rooted patterns, organ-system-rooted subgraphs suggest that intermediate physiological systems may act as upstream modulators shaping specific white-matter pathways.

For example, the metabolic node exhibits a highly focused downstream footprint, with a single direct edge to the cingulum within the cingulate gyrus, consistent with a targeted metabolic-to-limbic pathway. By contrast, the musculoskeletal node exhibits a markedly broader downstream topology, featuring a multi-step cascade through brainstem and cerebellar relay tracts (e.g., pontine crossing tract and middle cerebellar peduncle) and extending to major projection fibers such as the corticospinal tract. This structure suggests that musculoskeletal alterations may propagate through motor-related relay circuitry before reaching large-scale output pathways.

4 Discussion and Conclusion

In this study, we present an integrated, missingness-aware graph learning framework to explore brain-body-disease interactions from incomplete, high dimensional observational data. By combining organ-specific masked autoencoders, dimensionality alignment, and DAG-GNN, our approach constructs modality-aligned node representations and learns a unified directed graph over brain regions, organ systems, and disease phenotypes. Experiments on the UK Biobank

cohort demonstrate improved robustness to missing organ-level features and identify interpretable candidate pathways linking systemic organ functions to white-matter vulnerabilities. Specifically, the learned graph suggests transdiagnostic convergent vulnerabilities at core neural hubs, metabolism-related systemic cascades, and organ-specific neuromodulatory patterns.

However, these findings should be interpreted with caution. Because this study is based on observational UK Biobank data, the learned DAG represents candidate directed dependencies for hypothesis generation rather than confirmed biological causality. DAG-based structure learning relies on assumptions such as causal sufficiency, the Markov condition, and faithfulness, which may be affected by hidden confounding from medication, lifestyle, head motion, temporal differences, and selection effects. In addition, organ biomarkers and dMRI features were collected across different UKB instances, so the inferred edges should not be regarded as definitive evidence of temporal or interventional effects. Future longitudinal cohorts, Mendelian randomization, interventional studies, and external validation are needed to confirm these brain-body-disease hypotheses.

Acknowledgments. This work was supported by the National Key R&D Program of China (No. 2023YFF1204803) and the Key Project of Jiangsu Provincial Natural Science Fund (No. BK20253028). This research has been conducted using the UK Biobank Resource under Application Number 419012.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Abdi, H., Williams, L.J.: Principal component analysis. Wiley interdisciplinary reviews: computational statistics **2**(4), 433–459 (2010)
2. Alfaro-Almagro, F., Jenkinson, M., Bangerter, N.K., Andersson, J.L., Griffanti, L., Douaud, G., Sotiropoulos, S.N., Jbabdi, S., Hernandez-Fernandez, M., Vallee, E., et al.: Image processing and quality control for the first 10,000 brain imaging datasets from uk biobank. Neuroimage **166**, 400–424 (2018)
3. Alzoubi, A., Abunaser, R., Khassawneh, A., Alfaqih, M., Khasawneh, A., Abdo, N.: The bidirectional relationship between diabetes and depression: a literature review. Korean journal of family medicine **39**(3), 137 (2018)
4. Boquet-Pujadas, A., Anagnostakis, F., Yang, Z., Tian, Y.E., Duggan, M.R., Erus, G., Srinivasan, D., Joynes, C.M., Bai, W., et al.: Multi-organ ai endophenotypes chart the heterogeneity of brain, eye and heart pan-disease. Nature Mental Health pp. 1–28 (2026)
5. Chen, C., Chen, S., Ye, Z., Shi, X., Ma, T., Shardell, M.: The effect of alcohol intake on brain white matter microstructural integrity: A new causal inference framework for incomplete phenomic data. Journal of the American Statistical Association pp. 1–13 (2025)
6. Feuerriegel, S., Frauen, D., Melnychuk, V., Schweisthal, J., Hess, K., Curth, A., Bauer, S., Kilbertus, N., Kohane, I.S., van der Schaar, M.: Causal machine learning for predicting treatment outcomes. Nature Medicine **30**(4), 958–968 (2024)

7. Gao, Y., Taschler, B., Amin, N., van Duijn, C., Hunter, D.J., Topiwala, A., Littlejohns, T.J.: Association between neuroticism and brain-wide structural outcomes: Mediation by vascular and mental conditions. Psychological Medicine **55**, e343 (2025)
8. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16000–16009 (2022)
9. Jiang, Y., Zhao, B., Wang, X., Tang, B., Peng, H., Luo, Z., Shen, Y., Wang, Z., Jiang, Z., Wang, J., et al.: Ukb-mdrmf: a multi-disease risk and multimorbidity framework based on uk biobank data. Nature Communications **16**(1), 3767 (2025)
10. Kelly, S., Jahanshad, N., Zalesky, A., Kochunov, P., Agartz, I., Alloza, C., Andreassen, O., Arango, C., Banaj, N., Bouix, S., et al.: Widespread white matter microstructural differences in schizophrenia across 4322 individuals: results from the enigma schizophrenia dti working group. Molecular psychiatry **23**(5), 1261–1269 (2018)
11. Koshiyama, D., Fukunaga, M., Okada, N., Morita, K., Nemoto, K., Usui, K., Yamamori, H., Yasuda, Y., Fujimoto, M., Kudo, N., et al.: White matter microstructural alterations across four major psychiatric disorders: mega-analysis study in 2937 individuals. Molecular psychiatry **25**(4), 883–895 (2020)
12. Little, R.J., Rubin, D.B.: Statistical analysis with missing data. john wiley & sons. New York (2002)
13. Mayer, E.A., Tillisch, K., Gupta, A., et al.: Gut/brain axis and the microbiota. The Journal of clinical investigation **125**(3), 926–938 (2015)
14. Moor, M., Banerjee, O., Abad, Z.S.H., Krumholz, H.M., Leskovec, J., Topol, E.J., Rajpurkar, P.: Foundation models for generalist medical artificial intelligence. Nature **616**(7956), 259–265 (2023)
15. Mori, S., Oishi, K., Jiang, H., Jiang, L., Li, X., Akhter, K., Hua, K., Faria, A.V., Mahmood, A., Woods, R., et al.: Stereotaxic white matter atlas based on diffusion tensor imaging in an icbm template. Neuroimage **40**(2), 570–582 (2008)
16. Opel, N., Thalamuthu, A., Milaneschi, Y., Grotegerd, D., Flint, C., Leenings, R., Goltermann, J., Richter, M., Hahn, T., Woditsch, G., et al.: Brain structural abnormalities in obesity: relation to age, genetic risk, and common psychiatric disorders: Evidence through univariate and multivariate mega-analysis including 6420 participants from the enigma mdd working group. Molecular psychiatry **26**(9), 4839–4852 (2021)
17. Repple, J., Gruber, M., Mauritz, M., de Lange, S.C., Winter, N.R., Opel, N., Goltermann, J., Meinert, S., Grotegerd, D., Leehr, E.J., et al.: Shared and specific patterns of structural brain connectivity across affective and psychotic disorders. Biological Psychiatry **93**(2), 178–186 (2023)
18. Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., Downey, P., Elliott, P., Green, J., Landray, M., et al.: Uk biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. PLoS medicine **12**(3), e1001779 (2015)
19. Tian, Y.E., Cole, J.H., Bullmore, E.T., Zalesky, A.: Brain, lifestyle and environmental pathways linking physical and mental health. Nature Mental Health **2**(10), 1250–1261 (2024)
20. Tian, Y.E., Cropley, V., Maier, A.B., Lautenschlager, N.T., Breakspear, M., Zalesky, A.: Heterogeneous aging across multiple organ systems and prediction of chronic disease and mortality. Nature medicine **29**(5), 1221–1231 (2023)

21. Wang, X., Liu, J., Hu, S.S., Liu, Z., Lu, H., Liu, L., for the Alzheimer's Disease Neuroimaging Initiative: Hilama: High-dimensional multi-omics mediation analysis with latent confounding. BMC Medical Research Methodology **25**(1), 239 (2025)
22. Xie, C., Xiang, S., Zheng, Y., Shen, C., Li, Y., Cheng, W., Vaidya, N., Zhang, Z., Robinson, L., Winterer, J., et al.: Hierarchical neurocognitive model of externalizing and internalizing comorbidity. Nature Mental Health pp. 1–15 (2026)
23. Yu, Y., Chen, J., Gao, T., Yu, M.: Dag-gnn: Dag structure learning with graph neural networks. In: International conference on machine learning. pp. 7154–7163. PMLR (2019)
24. Zhang, R., Zhang, B., Shen, C., Sahakian, B.J., Li, Z., Zhang, W., Zhao, Y., Li, Y., Feng, J., Cheng, W.: Associations of dietary patterns with brain health from behavioral, neuroimaging, biochemical and genetic analyses. Nature Mental Health **2**(5), 535–552 (2024)
25. Zheng, X., Aragam, B., Ravikumar, P.K., Xing, E.P.: Dags with no tears: Continuous optimization for structure learning. Advances in neural information processing systems **31** (2018)