HiLoGraph: Hierarchical-Longitudinal Brain Network Representation Learning

Tong Chen[1], Minheng Chen[1], Jing Zhang[1], Chao Cao[1], Li Su[2,3], Tianming Liu[4], and Dajiang Zhu[1(✉)]

[1] Department of Computer Science and Engineering, University of Texas at Arlington, USA
{dajiang.zhu@uta.edu}
[2] Sheffield Institute for Translational Neuroscience, University of Sheffield, UK
[3] Department of Psychiatry, University of Cambridge, UK
[4] School of Computing, University of Georgia, USA

Abstract. Longitudinal brain Magnetic Resonance Imaging (MRI) enables direct observation of within-subject structural change, helping disentangle aging-related trajectories from inter-individual anatomical variation. Recent self-supervised methods have shown that encoding such longitudinal change yields representations useful for brain age prediction and disease detection. These methods, however, learn directly from volumetric images rather than from brain network representations, overlooking inter-regional connectivity patterns. Graph-based approaches, on the other hand, model brain networks but typically at a single parcellation level, using cross-sectional data, without distinguishing normal aging from pathological change. We therefore propose **HiLoGraph**, a self-supervised framework that encodes multi-scale brain networks derived from morphological features and structural connectivity into a shared latent space, where longitudinal constraints enforce temporal consistency across within-subject visits. Trained exclusively on healthy subjects, the learned latent space serves as a normative reference from which subject-level deviations can be quantified for downstream disease detection. Evaluated on 1,373 healthy subjects from the Aging Adult Brain Connectome (AABC) dataset and an independent clinical cohort of 156 subjects, HiLoGraph achieves an AUC of 0.769 for three-class classification of cognitively normal (CN), Alzheimer's disease (AD), and Lewy body dementia (LBD), and 0.731 for AD vs. LBD, outperforming single-scale and non-longitudinal ablations. Code is available at `https://github.com/tongchen2010/hilograph`.

Keywords: Brain Aging · Hierarchical Representation Learning · Normative Modeling · Lewy Body Dementia

1 Introduction

The human brain undergoes continuous structural change throughout aging, including progressive cortical thinning, regional volume loss, and alterations in

white-matter connectivity [22,25]. Although these changes vary substantially across individuals, they still follow broadly consistent population-level trajectories [3]. Learning meaningful representations of these normative aging patterns from healthy populations is therefore clinically valuable. Deviations from the expected course of healthy aging have been linked to cognitive decline, greater disease severity, and increased risk of conversion to dementia [8,15]. In this sense, representations of healthy aging provide a principled reference for identifying pathological change at the individual level, even without disease-specific labels during training [17,23].

Normative modeling builds a reference of healthy brain variation and detects pathology as a deviation from this reference [17]. Classical normative models typically fit each voxel or region independently using Gaussian Process regression [18] or hierarchical Bayesian regression [14], scaling to large populations without leveraging the brain's network structure. In parallel, graph neural networks have been applied to brain network analysis [16,13], representing the brain as a connectome and preserving inter-regional relationships that independent normative models overlook. Chen et al. [7] further extended this direction by integrating coarse atlas-level and fine-grained 3-hinge gyri (3HG) [5] morphological features, demonstrating that multi-scale fusion improves differentiation between AD and LBD. However, all the above methods learn from cross-sectional data and do not capture within-subject longitudinal change. Longitudinal self-supervised methods, such as Longitudinally-consistent Self-Organized Representation (LSOR) [19], Longitudinal Neighbourhood Embedding (LNE) [20], and Longitudinal Self-Supervised Learning (LSSL) [27] exploit within-subject repeated MRI data to capture genuine individual change, but operate on volumetric MRI rather than brain network representations and are often trained on mixed healthy and diseased cohorts, which can entangle disease-related variance with normative aging patterns. To our knowledge, no existing method jointly addresses multi-scale brain network encoding, longitudinal learning, and training on purely healthy data.

We propose HiLoGraph, a self-supervised framework that learns hierarchical brain network representations from longitudinal healthy data, jointly addressing the limitations of prior work. Specifically, our framework constructs two graphs to represent the subjects: an atlas-level region of interest (ROI) graph and a fine-grained landmarks-based graph (3HG) [5], each defined by morphological node features and structural connectivity (SC). The model consists of two components: a hierarchical graph encoder that fuses the two scales into a shared latent space through cross-scale message passing, and a top-down decoder that reconstructs ROI-level features from this latent space to enforce cross-scale consistency. For subjects with repeated visits, a trajectory smoothness loss and a contrastive ordering loss constrain within-subject latent trajectories to reflect genuine aging change. The resulting latent space, learned entirely from healthy subjects, provides a healthy reference representation from which subject-specific deviation features can be derived for downstream clinical analysis without retraining. We evaluate HiLoGraph on 1,373 healthy subjects from AABC [4] and

an in-house clinical cohort of 156 subjects for downstream disease classification. HiLoGraph achieves an AUC of 0.769 for three-class classification of CN, AD, and LBD, outperforming single-scale and non-longitudinal variants (permutation test $p<0.002$).

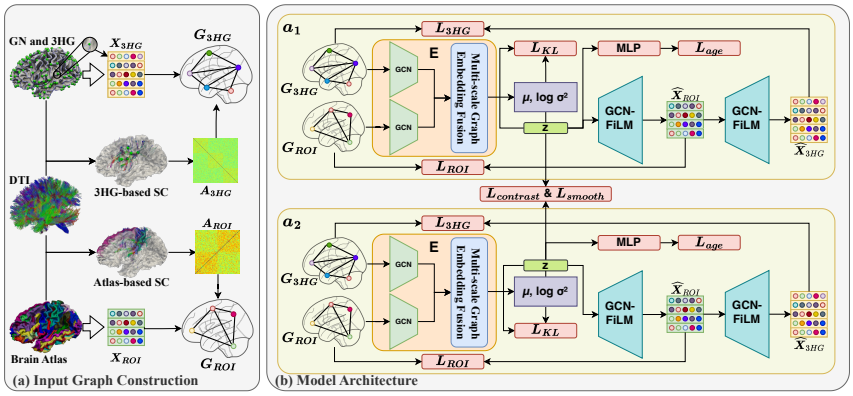

Fig. 1. Overview of our proposed framework. **(a)** Multi-scale brain graph construction: 3HG-based graph, $\mathcal{G}_{3HG}$, and an atlas-based graph, $\mathcal{G}_{ROI}$, constructed with morphological features and SC. **(b)** Model Architecture: the multi-scale graphs encoded with Hierarchical graph encoder E, and projected to a shared latent space $\mathbf{z}$ using the reparameterization trick, then reconstructed through top-down GCN-FiLM decoders. Longitudinal losses $\mathcal{L}_{\text{contrast}}$ and $\mathcal{L}_{\text{smooth}}$ enforce temporal consistency for subjects with two or more visits.

2 Materials and Method

2.1 Datasets and Preprocessing

Datasets. We used two datasets in this study. AABC Release 2, an extension of HCP-Aging [4], served as the normative training set. After quality control, 2,726 paired T1-weighted structural MRI and diffusion MRI data from 1,373 healthy adults (age 36–90 years; 1–4 visits per subject) were used. For downstream evaluation, we used a separate clinical cohort of 156 subjects with single-timepoint T1-weighted structural MRI and diffusion MRI, obtained from the University of Cambridge, with a subset from the Multimodal Imaging in Lewy Body Disorders (MILOS) project. The cohort included CN (n=20), AD spectrum (n=56), and LBD spectrum (n=80).

Preprocessing. The AABC dataset was released pre-processed following the HCP minimal preprocessing pipeline [11]. For the clinical cohort, T1-weighted images were denoised and bias-field-corrected [24], and cortical surfaces were

reconstructed with FreeSurfer [10] and resampled to 163,842 vertices per hemisphere. Diffusion MRI data were denoised and corrected for susceptibility distortions [1] and eddy currents [2] using FSL [12]. Diffusion and T1 volumes were then co-registered, and whole-brain deterministic tractography was performed with DSI Studio [26].

2.2 Model Architecture

An overview of the proposed framework is shown in Fig. 1. From each subject's T1 and diffusion MRI, we construct two graphs: a fine-grained 3HG graph and a coarse ROI graph. A hierarchical graph encoder fuses the two brain networks into a shared VAE latent space, and a top-down decoder reconstructs node features from coarse to fine, enforcing cross-scale consistency. For subjects with repeated visits, longitudinal losses constrain within-subject latent trajectories.

Graph construction. From each subject's T1 and diffusion MRI, we construct two graphs. For the landmark-based graph, 3HG landmarks are extracted on the white matter surface using the automated pipeline described in [5,6], with morphometric features computed from a 4-ring vertex neighbourhood around each landmark. The 3HG graph $\mathcal{G}_{\text{3HG}}=(\boldsymbol{X}_{\text{3HG}}, \boldsymbol{A}_{\text{3HG}})$ has $\sim$400 nodes, where $\boldsymbol{X}_{\text{3HG}}\in\mathbb{R}^{N\times 5}$ contains five cortical morphometric features (thickness, average convexity, mean curvature, surface area, volume) and $\boldsymbol{A}_{\text{3HG}}$ is SC with edge weights defined as log-normalised fiber counts from diffusion tractography. For the ROI graph $\mathcal{G}_{\text{ROI}}=(\boldsymbol{X}_{\text{ROI}}, \boldsymbol{A}_{\text{ROI}})$, the Destrieux atlas parcellation [9] from FreeSurfer was used and there are 148 nodes with the same morphometric features averaged within each parcel and ROI-level SC. A binary membership matrix $\boldsymbol{M}\in\{0,1\}^{N\times 148}$ assigns each 3HG node to its enclosing ROI.

Hierarchical graph encoder. A 3-layer GCN encodes the 3HG graph into node embeddings $\boldsymbol{H}_{\text{3HG}} \in \mathbb{R}^{N\times d}$. These are pooled to ROI resolution via the membership matrix:

$$\boldsymbol{H}_{\text{pool}} = \boldsymbol{D}^{-1}\boldsymbol{M}^{\top}\boldsymbol{H}_{\text{3HG}} \in \mathbb{R}^{148\times d} \tag{1}$$

where $\boldsymbol{D} = \text{diag}(\boldsymbol{M}^{\top})$ counts the number of 3HG nodes assigned to each ROI, and each row is divided by the number of assigned 3HG nodes. A 2-layer GCN independently encodes the ROI graph: $\boldsymbol{H}_{\text{ROI}}=\text{GCN}_{\text{ROI}}(\boldsymbol{X}_{\text{ROI}}, \boldsymbol{A}_{\text{ROI}})\in\mathbb{R}^{148\times d}$. The pooled 3HG embeddings are concatenated with their corresponding ROI embeddings, mean-pooled across nodes, and passed through a multilayer perceptron (MLP) to produce the VAE parameters:

$$\boldsymbol{\mu}, \log\boldsymbol{\sigma}^2 = \text{MLP}\big(\bar{\boldsymbol{h}}\big), \quad \boldsymbol{z} = \boldsymbol{\mu} + \boldsymbol{\epsilon}\odot\boldsymbol{\sigma}, \quad \boldsymbol{\epsilon}\sim\mathcal{N}(\boldsymbol{0}, \boldsymbol{I}) \tag{2}$$

where $\bar{\boldsymbol{h}} = \frac{1}{148}\sum_i[\boldsymbol{H}_{\text{pool},i}; \boldsymbol{H}_{\text{ROI},i}]$ is the mean-pooled concatenated representation.

Top-down GCN-FiLM decoders. The decoder reconstructs node features in a coarse-to-fine manner. First, the latent $\boldsymbol{z}$ is broadcast to all 148 ROI nodes, concatenated with learnable node-identity embeddings, and decoded by a GCN

with FiLM [21] conditioning to produce $\hat{\boldsymbol{X}}_{\mathrm{ROI}} \in \mathbb{R}^{148\times 5}$. For 3HG decoding, each 3HG node receives its parent ROI's decoded features via the membership matrix, $\boldsymbol{P} = \boldsymbol{M}\hat{\boldsymbol{X}}_{\mathrm{ROI}} \in \mathbb{R}^{N\times 5}$. This prior is concatenated with the broadcast $\boldsymbol{z}$ and decoded by a second GCN-FiLM to produce $\hat{\boldsymbol{X}}_{\mathrm{3HG}} \in \mathbb{R}^{N\times 5}$.

2.3 Training Objective

The training objective comprises three groups of losses: reconstruction, longitudinal, and age prediction.

Reconstruction. $\mathcal{L}_{\mathrm{ROI}} = \mathrm{MSE}(\hat{\boldsymbol{X}}_{\mathrm{ROI}}, \boldsymbol{X}_{\mathrm{ROI}})$ and $\mathcal{L}_{\mathrm{3HG}} = \mathrm{MSE}(\hat{\boldsymbol{X}}_{\mathrm{3HG}}, \boldsymbol{X}_{\mathrm{3HG}})$ supervise the ROI and 3HG decoders respectively. $\mathcal{L}_{\mathrm{KL}} = D_{\mathrm{KL}}(q(\boldsymbol{z}|\boldsymbol{x})\|p(\boldsymbol{z}))$ is the standard VAE prior, where $q(\boldsymbol{z}|\boldsymbol{x}) = \mathcal{N}(\boldsymbol{\mu}, \mathrm{diag}(\boldsymbol{\sigma}^2))$ and $p(\boldsymbol{z}) = \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$.

Longitudinal constraints. For subject i observed at ages $a_1 < a_2$ with embeddings $\boldsymbol{z}_{i,1}, \boldsymbol{z}_{i,2}$, a smoothness loss penalizes large changes in the latent space between visits:

$$\mathcal{L}_{\mathrm{smooth}} = \frac{1}{|\mathcal{P}|} \sum_{(i,a_1,a_2)\in\mathcal{P}} \left\| \frac{\boldsymbol{z}_{i,2} - \boldsymbol{z}_{i,1}}{a_2 - a_1} \right\|^2 \tag{3}$$

where $\mathcal{P}$ is the set of within-subject visit pairs. A contrastive loss pulls same-subject visits closer than different-subject visits, where $\mathrm{sim}(\cdot,\cdot)$ denotes cosine similarity and τ is a temperature parameter:

$$\mathcal{L}_{\mathrm{contrast}} = -\frac{1}{|\mathcal{P}|} \sum_{(i,j)\in\mathcal{P}} \log \frac{\exp(\mathrm{sim}(\boldsymbol{z}_i, \boldsymbol{z}_j)/\tau)}{\sum_{k\neq i} \exp(\mathrm{sim}(\boldsymbol{z}_i, \boldsymbol{z}_k)/\tau)} \tag{4}$$

These losses are computed only for subjects with at least two visits; single-visit subjects contribute to the reconstruction and KL terms only.

Age prediction. $\mathcal{L}_{\mathrm{age}} = \mathrm{MSE}(f_{\mathrm{age}}(\boldsymbol{z}), a)$, where f_{age} is a two-layer MLP that predicts chronological age from the latent vector $\boldsymbol{z}$. This auxiliary loss encourages age-related information to be encoded in the latent space.

Total loss. The overall objective function is:

$$\mathcal{L} = \lambda_r \mathcal{L}_{\mathrm{ROI}} + \lambda_f \mathcal{L}_{\mathrm{3HG}} + \lambda_{\mathrm{KL}} \mathcal{L}_{\mathrm{KL}} + \lambda_s \mathcal{L}_{\mathrm{smooth}} + \lambda_t \mathcal{L}_{\mathrm{contrast}} + \lambda_a \mathcal{L}_{\mathrm{age}}. \tag{5}$$

3 Experiments and Results

3.1 Experiments

Implementation details. The 3HG encoder consists of three GCN layers and the ROI encoder of two GCN layers, both with hidden dimension d=64, while the VAE bottleneck produces a 64-dimensional latent vector. Both the ROI and 3HG decoders use three GCN-FiLM [21] layers and 32-dimensional node identity embeddings. We train with Adam at learning rate 10^{-3} and weight decay 10^{-5} for 100 epochs with early stopping (patience 20), using 5-fold subject-level cross-validation. Loss weights were selected via grid search on a held-out validation set:

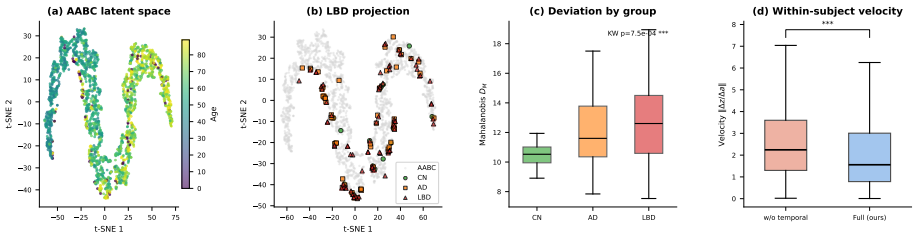

Fig. 2. Latent-space analysis of the full model. **(a)** t-SNE of AABC healthy latent codes coloured by age. **(b)** External clinical cohort projected into the same space (CN: green, AD: orange, LBD: red). **(c)** Mahalanobis deviation D_M by diagnostic group. **(d)** Within-subject latent velocity with and without temporal loss.

Table 1. Classification on the LBD dataset using deviation features **d**. Logistic regression with balanced class weights and repeated 5-fold stratified cross-validation.

	3-class		AD vs CN	AD vs LBD	CN vs LBD
Model	$\text{AUC}_{\text{ovr}}\uparrow$	BAcc↑	AUC↑	AUC↑	AUC↑
Baselines					
MLP	$.732 \pm .022$	$.566 \pm .036$	$.840 \pm .021$	$.679 \pm .031$	$.848 \pm .032$
GCN	$.689 \pm .019$	$.510 \pm .033$	$.805 \pm .035$	$.637 \pm .038$	$.807 \pm .022$
GAT	$.725 \pm .018$	$.547 \pm .028$	$.805 \pm .048$	$.674 \pm .027$	$.820 \pm .031$
LNE [20]	$.717 \pm .027$	$.531 \pm .036$	$.786 \pm .027$	$.676 \pm .021$	$.809 \pm .031$
LSSL [27]	$.758 \pm .013$	$.586 \pm .017$	$.837 \pm .031$	$\mathbf{.731 \pm .020}$	$.849 \pm .006$
Ablations					
3HG-only	$.739 \pm .015$	$.565 \pm .017$	$.796 \pm .047$	$.703 \pm .031$	$.847 \pm .014$
ROI-only	$.696 \pm .047$	$.517 \pm .067$	$.781 \pm .021$	$.644 \pm .059$	$.816 \pm .020$
Ours w/o long	$.689 \pm .025$	$.506 \pm .034$	$.794 \pm .041$	$.644 \pm .025$	$.785 \pm .030$
Ours	$\mathbf{.769 \pm .028}$	$\mathbf{.598 \pm .031}$	$\mathbf{.844 \pm .033}$	$\mathbf{.731 \pm .047}$	$\mathbf{.854 \pm .021}$

λ_r=1, λ_f=0.1, λ_{KL}=0.01, λ_s=0.5, λ_t=0.1, λ_a=0.1, with contrastive temperature τ=0.07.

Baselines. All baselines use the same training and evaluation protocol, with only the encoder architecture or input varied. MLP replaces graph message passing with an MLP applied to node features followed by mean pooling, thereby discarding graph structure. GCN and GAT use a 3-layer encoder on the 3HG graph, differing only in the message-passing operator. Two single-scale ablations isolate each branch of the hierarchy: 3HG-only removes the ROI encoder, whereas ROI-only removes the 3HG encoder. Finally, w/o long (in Table 1) and w/o temporal (in Fig. 3 and Fig. 2) use the full architecture but remove $\mathcal{L}_{\text{smooth}}$ and $\mathcal{L}_{\text{contrast}}$, isolating the contribution of the longitudinal constraints.

Normative deviation. After training, the encoder is frozen. Following [23], we fit age- and sex-conditioned baselines ($\mu_k(a, s)$, $\sigma_k(a, s)$) on the healthy training

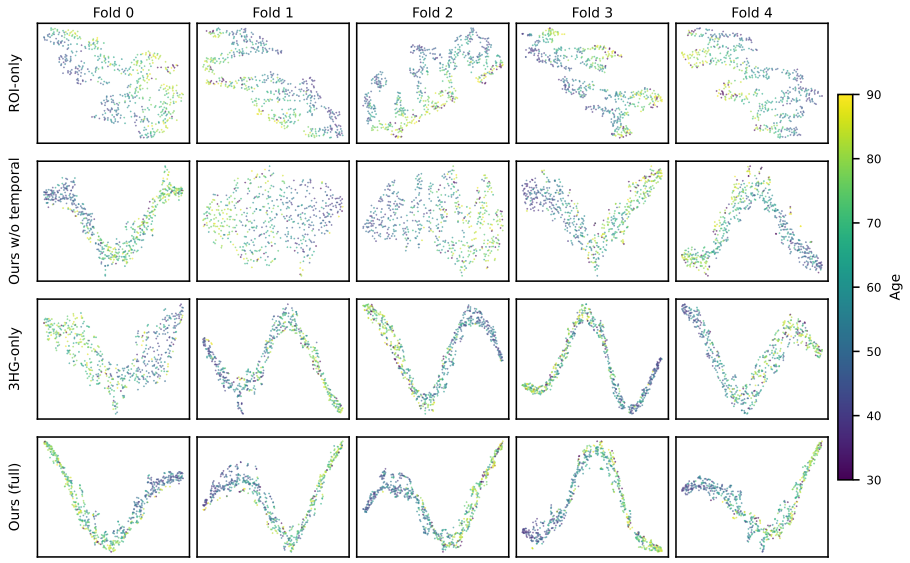

Fig. 3. t-SNE of AABC healthy latent codes coloured by age across four model variants (rows) and five cross-validation folds (columns). A smooth colour gradient indicates a well-organized aging manifold.

Table 2. Feature ablation on our full model. Deviation features consistently outperform raw latent codes.

	3-class		AD vs CN	AD vs LBD	CN vs LBD
Feature	AUC_{ovr}	BAcc	AUC	AUC	AUC
Raw $\mathbf{z}$	$.684 \pm .047$	$.493 \pm .053$	$.726 \pm .044$	$.673 \pm .042$	$.782 \pm .055$
Deviation $\mathbf{d}$	$\mathbf{.769 \pm .028}$	$\mathbf{.598 \pm .031}$	$.844 \pm .033$	$.731 \pm .047$	$\mathbf{.854 \pm .021}$
Combined $[\mathbf{z}; \mathbf{d}]$	$.767 \pm .041$	$.578 \pm .046$	$\mathbf{.856 \pm .035}$	$\mathbf{.734 \pm .050}$	$.854 \pm .025$

latent space using robust quadratic regression. Subject-specific deviation scores $d_k{=}(z_k{-}\mu_k)/\sigma_k$ serve as features for a downstream logistic regression classifier.

3.2 Results

Latent space captures aging and disease structure. The learned latent space shows a clear age gradient across the healthy AABC cohort (Fig. 2a), shaped by the auxiliary age-prediction loss and longitudinal constraints. When the LBD dataset is projected into this normative space without fine-tuning (Fig. 2b), CN subjects largely overlap with the healthy distribution, whereas AD and LBD subjects are shifted toward older apparent brain age. Age-adjusted Mahalanobis deviation scores reveal a significant group effect (Kruskal–Wallis $p{=}7.5{\times}10^{-4}$; Fig. 2c), with LBD and AD patients showing greater deviation

from healthy norms than CN subjects, despite the encoder being trained exclusively on healthy data.

Latent space organization across ablations. Fig. 3 shows t-SNE projections of the healthy AABC latent space across all five folds for four model variants. Only the full model consistently produces a smooth age gradient in every fold. The 3HG-only ablation captures age structure in most folds but fragments in fold 0, while removing longitudinal constraints disrupts age ordering in two of five folds. The ROI-only ablation fails to recover any age structure, indicating that the fine-grained 3HG graph carries the primary aging information. These results suggest that both hierarchical fusion and longitudinal constraints are necessary to reliably encode aging in the latent space.

Classification results. Table 1 reports results on the external LBD cohort. Our full model achieves a 3-class AUC of 0.769, with binary classification AUCs of 0.844 for AD vs. CN, 0.854 for CN vs. LBD, and 0.731 for AD vs. LBD. The classification between AD and LBD is the most challenging, as both groups exhibit neurodegeneration and the model must distinguish the pattern rather than the presence of deviation. Removing the longitudinal constraints degrades AUC from 0.769 to 0.689 for 3-class classification and from 0.731 to 0.644 for AD vs. LBD, confirming that learning smoother healthy aging trajectories improves downstream neurodegenerative disease classification. Incorporating longitudinal constraints also yields significantly smoother within-subject trajectories ($p<0.001$, Mann-Whitney U, Fig. 2d). The hierarchical encoder outperforms the single-scale GCN by 0.094 AUC on AD vs. LBD.

Effect of normative deviation. Table 2 compares raw latent vectors $\mathbf{z}$, normative deviations $\mathbf{d}$, and their concatenations. Deviation features improve the 3-class AUC by 0.085 over raw vectors, with the largest improvement observed on AD vs. CN by 0.118, where age-corrected deviation clearly separates pathological decline from normal aging. Concatenations $[\mathbf{z}; \mathbf{d}]$ provide little additional benefit, indicating that the deviations capture most of the disease-relevant pattern.

Statistical significance. A permutation test (1,000 label shuffles) confirms that the full model significantly exceeds chance (empirical $p<0.002$ in every fold; null AUC 0.500 ± 0.051).

4 Conclusion

We presented HiLoGraph, a framework that learns brain network representations exclusively from healthy longitudinal data. Multi-scale fusion and longitudinal constraints each contribute meaningfully, jointly achieving AUC 0.769 (CN vs. AD vs. LBD) and 0.731 (AD vs. LBD) on an independent clinical cohort. These results are achieved using a frozen encoder without any disease-specific fine-tuning, suggesting that structurally rich representations of healthy aging can serve as a general-purpose foundation for detecting diverse neurodegenerative conditions. The current evaluation is limited by the modest clinical cohort and the absence of explicit cross-site harmonization. Future work will address

these through larger multi-site validation with harmonization, region-level interpretability, and incorporation of functional connectivity.

Acknowledgments. This study was supported by the National Institutes of Health (R01AG075582 and R01NS128534) and funded by the Sir Jules Thorn Charitable Trust (05/JTA), Wellcome Trust Intermediate Clinical Fellowship (WT088441MA), Alzheimer's Research UK Senior Research Fellowship (ARUK-SRF2017B-1), The Lewy body Society (LBS/002/2019) and NIHR Biomedical Research Centres in Sheffield, Cambridge and Newcastle. We also thank Andrew M. Blamire, John-Paul Taylor, Li Su and John T. O'Brien for sharing the data.

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

References

1. Andersson, J.L., Skare, S., Ashburner, J.: How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. Neuroimage **20**(2), 870–888 (2003)
2. Andersson, J.L., Sotiropoulos, S.N.: An integrated approach to correction for off-resonance effects and subject movement in diffusion mr imaging. Neuroimage **125**, 1063–1078 (2016)
3. Bethlehem, R.A., Seidlitz, J., White, S.R., Vogel, J.W., Anderson, K.M., Adamson, C., Adler, S., Alexopoulos, G.S., Anagnostou, E., Areces-Gonzalez, A., et al.: Brain charts for the human lifespan. Nature **604**(7906), 525–533 (2022)
4. Bookheimer, S.Y., Salat, D.H., Terpstra, M., Ances, B.M., Barch, D.M., Buckner, R.L., Burgess, G.C., Curtiss, S.W., Diaz-Santos, M., Elam, J.S., et al.: The lifespan human connectome project in aging: an overview. Neuroimage **185**, 335–348 (2019)
5. Cao, C., Chen, T., Zhao, N., Chen, M., Qu, M., Zhang, Z., Shi, X., Li, X., Liu, T., Zhang, L.: Gyral-sulcal-net: an integrated network representation of brain folding patterns. In: 2026 IEEE 23rd International Symposium on Biomedical Imaging (ISBI). pp. 1–5 (2026). https://doi.org/10.1109/ISBI61048.2026.11515619
6. Chen, H., Li, Y., Ge, F., Li, G., Shen, D., Liu, T.: Gyral net: A new representation of cortical folding organization. Medical image analysis **42**, 14–25 (2017)
7. Chen, T., Chen, M., Zhang, J., Zhuang, Y., Cao, C., Yu, X., Lyu, Y., Zhang, L., Su, L., Liu, T., et al.: A unified continuous staging framework for alzheimer's disease and lewy body dementia via hierarchical anatomical features. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 13–23. Springer (2025)
8. Cole, J.H., Poudel, R.P., Tsagkrasoulis, D., Caan, M.W., Steves, C., Spector, T.D., Montana, G.: Predicting brain age with deep learning from raw imaging data results in a reliable and heritable biomarker. NeuroImage **163**, 115–124 (2017)
9. Destrieux, C., Fischl, B., Dale, A., Halgren, E.: Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature. Neuroimage **53**(1), 1–15 (2010)
10. Fischl, B.: Freesurfer. Neuroimage **62**(2), 774–781 (2012)
11. Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., et al.: The minimal preprocessing pipelines for the human connectome project. Neuroimage **80**, 105–124 (2013)

12. Jenkinson, M., Beckmann, C.F., Behrens, T.E., Woolrich, M.W., Smith, S.M.: Fsl. Neuroimage **62**(2), 782–790 (2012)
13. Jiang, H., Cao, P., Xu, M., Yang, J., Zaiane, O.: Hi-gcn: A hierarchical graph convolution network for graph embedding learning of brain network and brain disorders prediction. Computers in Biology and Medicine **127**, 104096 (2020)
14. Kia, S.M., Huijsdens, H., Rutherford, S., de Boer, A., Dinga, R., Wolfers, T., Berthet, P., Mennes, M., Andreassen, O.A., Westlye, L.T., et al.: Closing the lifecycle of normative modeling using federated hierarchical bayesian regression. Plos one **17**(12), e0278776 (2022)
15. Leonardsen, E.H., Peng, H., Kaufmann, T., Agartz, I., Andreassen, O.A., Celius, E.G., Espeseth, T., Harbo, H.F., Høgestøl, E.A., De Lange, A.M., et al.: Deep neural networks learn general and clinically relevant representations of the ageing brain. NeuroImage **256**, 119210 (2022)
16. Li, X., Zhou, Y., Dvornek, N., Zhang, M., Gao, S., Zhuang, J., Scheinost, D., Staib, L.H., Ventola, P., Duncan, J.S.: Braingnn: Interpretable brain graph neural network for fmri analysis. Medical image analysis **74**, 102233 (2021)
17. Marquand, A.F., Kia, S.M., Zabihi, M., Wolfers, T., Buitelaar, J.K., Beckmann, C.F.: Conceptualizing mental disorders as deviations from normative functioning. Molecular psychiatry **24**(10), 1415–1424 (2019)
18. Marquand, A.F., Wolfers, T., Mennes, M., Buitelaar, J., Beckmann, C.F.: Beyond lumping and splitting: a review of computational approaches for stratifying psychiatric disorders. Biological psychiatry: cognitive neuroscience and neuroimaging **1**(5), 433–447 (2016)
19. Ouyang, J., Zhao, Q., Adeli, E., Peng, W., Zaharchuk, G., Pohl, K.M.: Lsor: Longitudinally-consistent self-organized representation learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 279–289. Springer (2023)
20. Ouyang, J., Zhao, Q., Adeli, E., Sullivan, E.V., Pfefferbaum, A., Zaharchuk, G., Pohl, K.M.: Self-supervised longitudinal neighbourhood embedding. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 80–89. Springer (2021)
21. Perez, E., Strub, F., De Vries, H., Dumoulin, V., Courville, A.: Film: Visual reasoning with a general conditioning layer. In: Proceedings of the AAAI conference on artificial intelligence. vol. 32 (2018)
22. Raz, N., Lindenberger, U., Rodrigue, K.M., Kennedy, K.M., Head, D., Williamson, A., Dahle, C., Gerstorf, D., Acker, J.D.: Regional brain changes in aging healthy adults: general trends, individual differences and modifiers. Cerebral cortex **15**(11), 1676–1689 (2005)
23. Rutherford, S., Fraza, C., Dinga, R., Kia, S.M., Wolfers, T., Zabihi, M., Berthet, P., Worker, A., Verdi, S., Andrews, D., et al.: Charting brain growth and aging at high spatial precision. elife **11**, e72904 (2022)
24. Tustison, N.J., Avants, B.B., Cook, P.A., Zheng, Y., Egan, A., Yushkevich, P.A., Gee, J.C.: N4itk: improved n3 bias correction. IEEE transactions on medical imaging **29**(6), 1310–1320 (2010)
25. Verdi, S., Kia, S.M., Yong, K.X., Tosun, D., Schott, J.M., Marquand, A.F., Cole, J.H.: Revealing individual neuroanatomical heterogeneity in alzheimer disease using neuroanatomical normative modeling. Neurology **100**(24), e2442–e2453 (2023)
26. Yeh, F.C., Verstynen, T.D., Wang, Y., Fernández-Miranda, J.C., Tseng, W.Y.I.: Deterministic diffusion fiber tracking improved by quantitative anisotropy. PloS one **8**(11), e80713 (2013)

27. Zhao, Q., Liu, Z., Adeli, E., Pohl, K.M.: Longitudinal self-supervised learning. Medical image analysis **71**, 102051 (2021)