POLITECNICO DI TORINO
Repository ISTITUZIONALE

Overcoming Modality Scarcity in Parkinson's Diagnosis Via Generative Geometric Deep Learning

Original
Overcoming Modality Scarcity in Parkinson's Diagnosis Via Generative Geometric Deep Learning / Barbieri, E., Monaco, S., Apiletti, D.. - (2026), pp. 289-294. (2026 IEEE International Workshop on Metrology for Living Environment, MetroLivEnv 2026 Cambridge (UK) 14-16 July 2026) [10.1109/metrolivenv70468.2026.11659857].

Availability:
This version is available at: 11583/3018338 since: 2026-11-23T15:47:02Z

Publisher:
Institute of Electrical and Electronics Engineers

Published
DOI:10.1109/metrolivenv70468.2026.11659857

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright
IEEE postprint/Author's Accepted Manuscript

©2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

28 November 2026

Overcoming Modality Scarcity in Parkinson's Diagnosis via Generative Geometric Deep Learning

Emanuele Barbieri, Simone Monaco, and Daniele Apiletti
Dept. of Control and Computing Engineering
Politecnico di Torino
Turin, Italy
{name.surname}@polito.it

***Abstract*—Early detection of Parkinson's disease (PD) and accurate modeling of its progression remain challenging, particularly when attempting to fully exploit the complementary information provided by diverse neuroimaging modalities. In this work we propose GeNeuro, a geometric deep learning framework that integrates functional magnetic resonance imaging (fMRI) and diffusion tensor imaging (DTI) within a contrastive learning pipeline, enabling effective information extraction from both modalities while maintaining robustness in the absence of either.**

Our approach jointly models these modalities to improve both diagnostic and prognostic performance. Specifically, it achieves $\sim 10\%$ improvement in PD diagnosis accuracy compared to SOTA models while classifying PD, healthy controls and prodromal patients utilizing solely fMRI and DTI based neuroimaging in the specific circumstance where fMRI data is exceptionally scarce. The achieved results suggests low reliance on fMRI and the ability to reconstruct missing fMRI embeddings from existing DTI.

Beyond predictive performance, the proposed framework incorporates an interpretable mechanism that links model outputs to relevant brain substructures, potentially spanning multiple structurally or functionally connected regions. This enables the identification of disease-related patterns in brain organization.

***Index Terms*—Parkinson's Disease, Neuroimaging, Diagnosis, Severity, Progression**

I. Introduction

Parkinson's Disease (PD) is a rapidly growing neurodegenerative disorder characterized by the progressive death of dopaminergic neurons in the Substantia Nigra Pars Compacta [1]. Diagnosing PD at its early stages is critical to administer targeted treatment. While neuroimaging and deep learning show significant promise, current state-of-the-art models often struggle to effectively fuse multiple neuroimaging modalities or compensate for the absence of specific scans in patient data. Furthermore, existing diagnostic methods that rely solely on neuroimaging often fail to account for the prodromal class and reduce the classification problem to separating patients into only the PD and Healthy Controls classes.

To address these limitations, this work proposes a novel geometric deep learning framework that integrates two neuroimaging modalities commonly used in Parkinson's disease research: functional magnetic resonance imaging (fMRI), which captures brain activity, and diffusion tensor imaging (DTI), which characterizes white-matter structural connectivity. By jointly leveraging these complementary sources of information, the framework can more effectively model the complex neural alterations associated with the disease. In addition, a contrastive learning strategy is introduced to enable reliable predictions even when one imaging modality is unavailable, thereby improving the robustness and practical applicability of the approach.

The outcomes of this ongoing work are as follows:

- We develop a preprocessing pipeline that transforms fMRI and DTI acquisitions into weighted graph representations, enabling the extraction of meaningful functional and structural connectivity features.
- We train a graph neural network using a contrastive learning strategy to effectively integrate multimodal data and maintain robustness in the presence of missing modalities.
- We introduce an interpretability stage that links model predictions to relevant brain regions and substructures, ensuring the clinical and biological relevance of the results.

II. Related Works

Graph-based representations have emerged as a natural choice for modeling brain data, given the structural analogy between brain networks and graphs, where regions of interest (ROIs) correspond to nodes and their interactions to edges. Several works adopt this paradigm for multimodal integration. For instance, *IBrainGNN* [2] follows a conventional strategy in which ROIs derived from MRI parcellation serve as graph nodes, while edge weights encode structural and functional connectivity measures computed from DTI and fMRI, respectively. Similarly, *BrainMAP* [3] proposes a more refined fusion mechanism by first selecting disease-relevant regions

Data used in the preparation of this paper were obtained from the Parkinson's Progression Markers Initiative (PPMI) database (www.ppmi-info.org/data). For up-to-date information on the study, visit www.ppmi-info.org.

PPMI – a public-private partnership – is funded by the Michael J. Fox Foundation for Parkinson's Research and funding partners, including: 4D Pharma, AbbVie, AcureX, Allergan, Amathus Therapeutics, ASAP, AskBio, Avid Radiopharmaceuticals, BIAL, BioArctic, Biogen, Biohaven, BioLegend, BlueRock Therapeutics, Bristol Myers Squibb, Calico Labs, Capsida, Eli Lilly, Gain Therapeutics, GE HealthCare, Genentech, GSK, Golub Capital, Handl Therapeutics, Insitro, Jazz Pharmaceuticals, Johnson & Johnson, Lundbeck, Merck, Meso Scale Discovery, Mission Therapeutics, Neurocrine Biosciences, Neuron23, Neuropore, Pfizer, Piramal, Prevail Therapeutics, Roche, Sanofi, Servier, Takeda, Teva, UCB, Vanqua Bio, Verily, and Voyager Therapeutics.

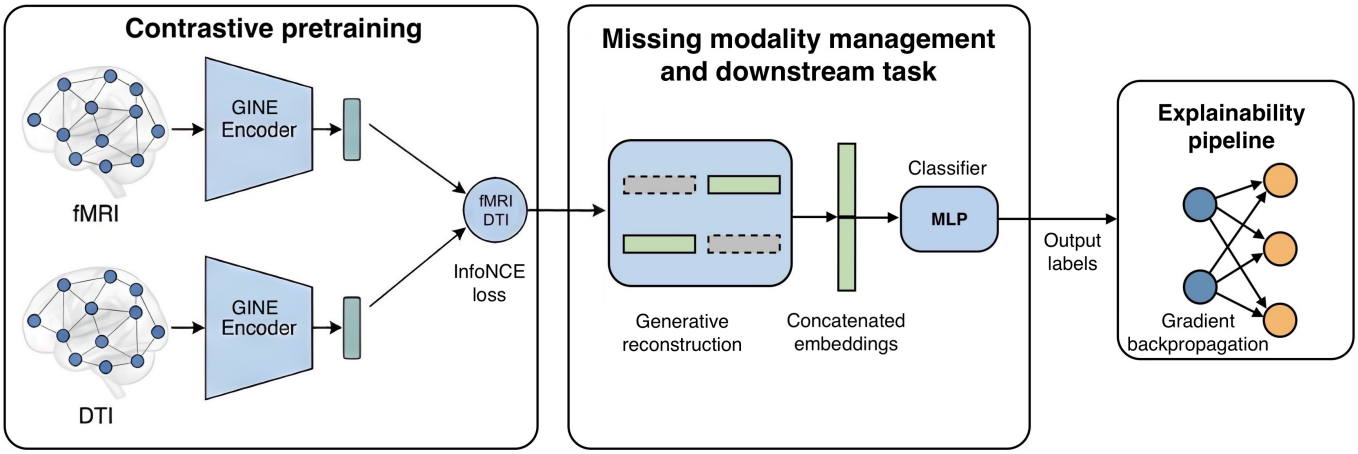


Fig. 1. Visual representation of the developed architecture

and then projecting structural and functional graphs into a shared attention space, enabling cross-modal alignment.

Despite their effectiveness, these approaches often rely on precise alignment between modalities, which can be difficult to achieve in practice. This limitation becomes more pronounced when incorporating modalities with differing spatial or temporal resolutions. For example, integrating other modalities, such as SPECT data, which has low temporal resolution, poses challenges for accurate connectivity estimation. To address this, *NeuroMark* [4] introduces an fMRI-guided framework for constructing SPECT-based brain networks using spatially constrained independent component analysis. While this method facilitates cross-modal guidance, aligning the resulting SPECT representations with their fMRI counterparts remains complex, complicating their joint use within a unified framework.

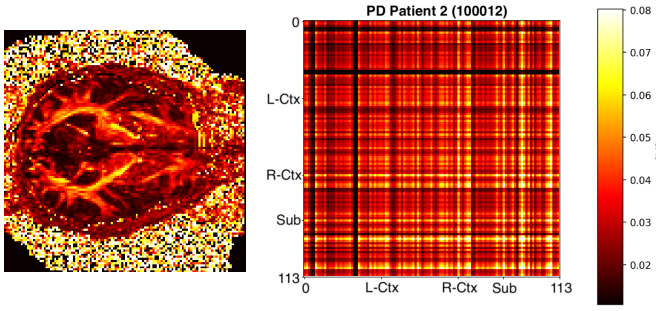
In contrast to prior work, our approach leverages a graph-based preprocessing pipeline combined with contrastive learning to integrate multimodal data flexibly. This design reduces the need for strict alignment across different modalities while improving robustness to missing data. Our work focuses on the analysis of DTI and fMRI modalities, but can be extended to others. Additionally, by incorporating an interpretability stage, our method bridges the gap between predictive performance and neuroscientific insight, linking model outputs to meaningful brain regions and structures.

III. Methodology

A. Data Preprocessing and Brain Network Construction

The framework preprocesses raw neuroimaging data to produce robust graph structures, where nodes represent Regions of Interest (ROIs) and edges denote structural or functional connectivity. Namely, fMRI preprocessing produces Functional Connectivity (FC) matrices, which are constructed by extracting time-series data and calculating the Pearson correlation of BOLD signals between ROIs [5]. Then, DTI is used to produce structural connectivity (SC) matrices, which were derived by solving for the diffusion tensor D in the Stejskal-Tanner equation to track white matter streamlines, and computing the product of Fractional Anisotropy (FA) between ROIs.

Figure 2 shows an example of a raw DTI acquisition after a few preprocessing steps, alongside its resulting model input.



(a) (b)

Fig. 2. Example of raw DTI acquisition (left) after tensor fitting to solve the Stejskal-Tanner equation and corresponding structural connectivity matrix (right) built during data preprocessing. There are 113 regions in total, with each pair of regions associated with a value representing the strength of the structural connection between them.

B. The dataset

From Table I we may observe that our final post processed dataset has over 3000 unique patient-visit pairs balanced across fMRI and DTI. Out of these pairs only 604 appear in both modalities (signifying that most patient-vist pairs have missing modalities) and there is significant class imbalance, with no prodromal patient having a DTI scan and very few control patients having an fMRI scan.

Figure 3 highlights the distributions of edge weights in DTI and fMRI graphs. The first distribution reflects SC matrices that are known to be sparse and strictly positive since a negative connection would have no physical meaning when measuring white matter tract integrity. FC matrices have

TABLE I
DISTRIBUTION AND INTERSECTION OF (PATIENT, VISIT) PAIRS FOR FMRI AND DTI

Modality	Total Pairs	Group Breakdown			Intersection
		PD	Control	Prodromal	
fMRI	1544	727	66	751	604
DTI	1526	228	1298	0	

stronger connections, and BOLD signals in different ROIs that can be negatively correlated.

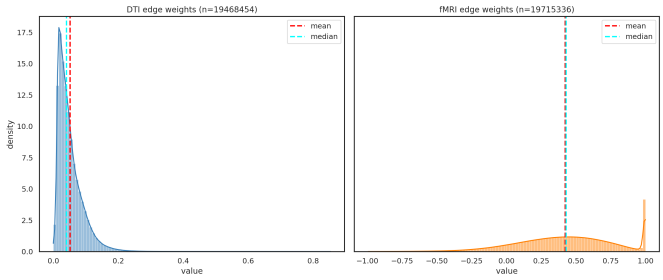


Fig. 3. Distribution of edge weights in DTI acquisitions (left) and fMRI acquisitions (right)

C. Model Architecture

The core training and feature extraction are summarized visually in Figure 1.

1) Contrastive pretraining: fMRI and DTI processed graphs corresponding to the same patient and visit are first aligned in a shared representation space through a bidirectional contrastive objective based on the InfoNCE loss [6]. To this end, each modality is processed by a dedicated graph encoder built upon GIN layers. We opted for this backbone for its superior ability to capture complex structural patterns, which make it well-suited for graph-classification tasks [7], [8]. Starting from brain graphs defined over a common set of atlas-derived regions, the two encoders generate modality-specific graph representations and compress them into fixed-length latent vectors of dimension 1024.

Because both structural and functional graphs are defined over the same node set, the resulting embeddings are inherently compatible and can be directly combined at the graph level. This shared anatomical support simplifies multimodal integration and enables the downstream model to jointly reason over structural and functional information without requiring additional graph-matching or alignment procedures.

The contrastive task is then defined: embeddings extracted from the fMRI and DTI acquisitions of the same subject at the same visit are treated as positive pairs, whereas embeddings from other subjects within the same mini-batch serve as negative samples. A symmetric InfoNCE formulation is adopted, in which each modality is alternately treated as the anchor and the target. This objective promotes high similarity between matched structural and functional representations while discouraging agreement between unrelated pairs, thereby inducing a latent space that captures cross-modal consistency at the subject level.

2) Missing modality management and downstream task: Clinical datasets often lack complete modality coverage, as patients may undergo only a subset of the required examinations. To address this limitation, a masked variational autoencoder (VAE) based on a Transformer architecture is employed, in which missing modalities are represented via learnable mask tokens. The resulting sequence is processed by the Transformer, and the reparameterization trick is applied to sample a global latent representation. This unified latent vector is then decoded to reconstruct high-fidelity embeddings for the missing patient data, enabling the model to operate with incomplete samples that would otherwise be inaccessible.

Following the pretraining phase, the fused multimodal representations are fed into a classification head and optimized for the downstream task. Specifically, the model performs clinical classification by assigning each subject to one of three categories: Parkinson's disease (PD), healthy control (HC), or prodromal.

D. Explainability Pipeline

To enhance the interpretability of the framework, we incorporate a unified gradient-based feature attribution module. This component quantifies the contribution of input graph features to the predictions by backpropagating gradients from the classification loss to the input space. Given a trained model, gradients are computed with respect to the input graph representations, enabling the estimation of feature-level sensitivity for both structural and functional connectivity patterns.

From these gradients, multiple attribution measures are derived by analyzing the absolute saliency of gradient magnitudes. To identify consistent and clinically meaningful biomarkers, individual attribution maps are aggregated across subjects. Relative saliency scores are computed by normalizing and averaging attribution values within each diagnostic group, thereby highlighting stable patterns that are robust to inter-subject variability.

By linking model predictions back to anatomically defined brain regions and connections, this interpretability module provides actionable insights into the neurobiological mechanisms captured by the model, facilitating its potential use in clinical decision support.

IV. RESULTS

The framework combines two distinct neuroimaging modalities from the Parkinson's Progression Markers Initiative (PPMI) dataset [9]. By mapping these modalities using Graph Neural Networks (GNNs), contrastive learning, and generative AI, the framework constructs a comprehensive shared latent space, imputes missing modalities, and successfully diagnoses PD while offering gradient-based explainability. Extensive Bayesian hyperparameter optimization was conducted to stabilize the architecture. A critical hyperparameter was the threshold for the sparsification of DTI and fMRI graphs, which

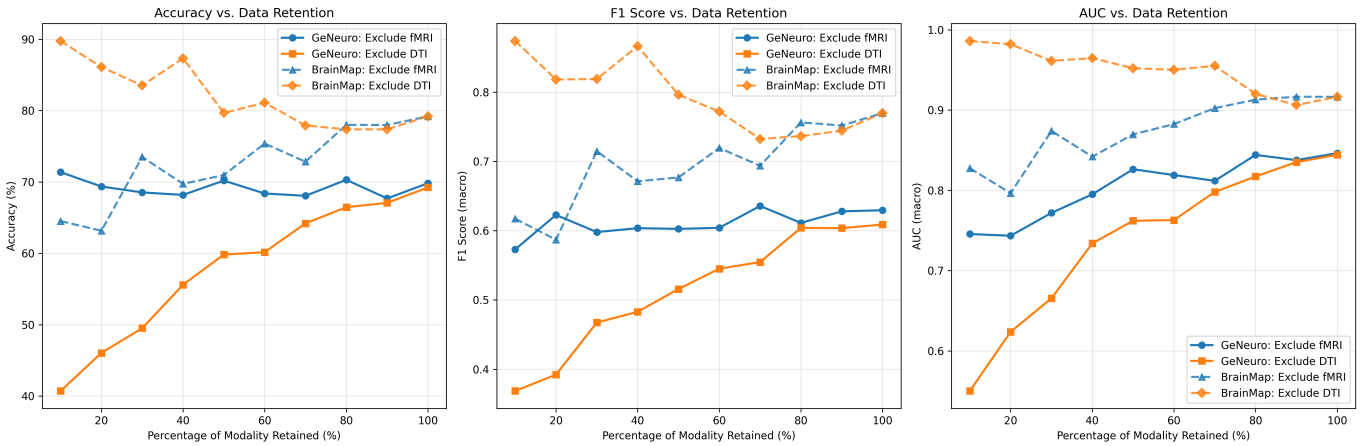


Fig. 4. Performance of *BrainMap* (dashed line) and GeNeuro (solid line) in the classification task reported in terms of accuracy, F1 score, and AUC metrics plotted against the varying of the percentage of samples of a certain modality that are retained.

converged to 0.60; effectively pruning all normalized edges below this value to retain only the strongest connections. All model stages were trained for 100 epochs using an AdamW optimizer, though learning rates varied significantly by phase. Specifically, contrastive pre-training utilized a rate of 0.00015, whereas generative reconstruction required a much lower rate of 0.000012 to maintain stability, and downstream tasks converged optimally at 0.01. Finally, the capacity of the model was defined by an embedding dimension of 1024, which was learned during pre-training and held fixed across the generative and downstream phases.

A. Predictive Performance

Table II reports the downstream task performance, with mean and standard deviation, when training the model using a patient-level stratified 5-fold cross-validation schema.

TABLE II
5-FOLD CROSS-VALIDATION PERFORMANCE ACROSS DOWNSTREAM TASKS ($n = 5$). RESULTS ARE REPORTED AS MEAN ± STANDARD DEVIATION.

Task & Metric	Score (Mean ± SD)
Clinical Classification	
Balanced Accuracy	0.7632 ± 0.0167
Macro F1-Score	0.6584 ± 0.0157
Macro AUC	0.8823 ± 0.0041

We designed an experiment to compare GeNeuro against state-of-the-art alternatives in terms of robustness to data scarcity for one of the two modalities. Figure 4 illustrates the sensitivity of BrainMap and our model to modality-specific data retention, measured by accuracy, F1-score, and AUC. Specifically, the x-axis of each plot represents the ratio of dataset samples retained for the reduced modality, while the other modality remains unvaried. In BrainMap, missing modalities are handled using zeroed adjacency matrices, whereas our model naturally addresses them during generation.

For BrainMap, performance shows a clear asymmetry in modality removal. As fMRI data are progressively excluded (dashed blue curves), all three metrics degrade substantially at low retention levels and recover only when a large proportion of fMRI samples is preserved. In contrast, excluding DTI data (dashed orange curves) results in comparatively strong performance, particularly at lower retention rates, where BrainMap achieves its highest scores. This counterintuitive trend suggests that BrainMap may not effectively leverage the complementarity of multimodal data. Instead, the inclusion of DTI data appears to dilute the contribution of fMRI, ultimately limiting the model's expressive capacity. The relatively higher performance when most DTI samples are removed supports the hypothesis that its fusion strategy may introduce noise or redundancy rather than useful signal.

Our model (solid curves) shows markedly different, more stable behavior. Although absolute performance is generally lower than BrainMap at high retention levels, the model benefits consistently from the gradual inclusion of data from both modalities. When fMRI data are excluded, performance remains relatively stable across retention levels, indicating reduced dependence on a single modality. More importantly, when DTI data are progressively included, all metrics improve steadily, without the degradation observed in BrainMap. This trend demonstrates a more effective integration of multimodal information, where additional data enhances rather than disrupts learned representations.

Overall, these results suggest that BrainMap is highly sensitive to modality imbalance and may suffer from suboptimal fusion, whereas our model, despite lower peak performance, demonstrates more robust and scalable multimodal learning, consistently benefiting from increased data availability across both modalities.

B. Biomarker Discovery and Clinical Insights

The explainability module successfully validated known PD pathology and provided nuanced biological insights by

TABLE III
TOP INTER-REGIONAL GLOBAL CONNECTIONS (EDGES) FOR 3-WAY CLASSIFICATION

Modality	Connection (Edge)	Relative Saliency	PD	HC	Prod.
fMRI	optic_chiasm ↔ rh_caudalmiddlefrontal	97.3x	-1.5×10^{-5}	-6.0×10^{-6}	-9.1×10^{-5}
fMRI	optic_chiasm ↔ rh_posteriorcingulate	83.9x	-1.8×10^{-5}	-4.0×10^{-6}	-5.4×10^{-5}
fMRI	optic_chiasm ↔ rh_superiorfrontal	74.8x	-1.5×10^{-5}	-5.0×10^{-6}	-8.5×10^{-5}
fMRI	optic_chiasm ↔ rh_caudalanteriorcingulate	69.8x	-9.0×10^{-6}	-8.0×10^{-6}	-7.8×10^{-5}
fMRI	optic_chiasm ↔ rh_precentral	63.3x	-1.2×10^{-5}	-8.0×10^{-6}	-3.9×10^{-5}
fMRI	cc_mid_anterior ↔ rh_posteriorcingulate	54.3x	-7.0×10^{-6}	-2.0×10^{-6}	-4.6×10^{-5}
fMRI	cc_mid_anterior ↔ rh_caudalmiddlefrontal	53.4x	-8.0×10^{-6}	-3.0×10^{-6}	-6.7×10^{-5}
fMRI	left_caudate ↔ optic_chiasm	50.0x	-1.1×10^{-5}	-2.0×10^{-6}	-3.0×10^{-5}
fMRI	optic_chiasm ↔ rh_postcentral	47.9x	-6.0×10^{-6}	-4.0×10^{-6}	-3.3×10^{-5}
fMRI	cc_mid_anterior ↔ rh_superiorfrontal	44.0x	-5.0×10^{-6}	-2.0×10^{-6}	-5.4×10^{-5}
DTI	rh_bankssts ↔ rh_caudalanteriorcingulate	486.9x	-1.6×10^{-5}	-1.1×10^{-5}	0.00
DTI	rh_caudalanteriorcingulate ↔ rh_middletemporal	451.6x	-1.8×10^{-5}	-1.2×10^{-5}	0.00
DTI	rh_caudalanteriorcingulate ↔ rh_inferiortemporal	425.0x	-1.9×10^{-5}	-1.4×10^{-5}	0.00
DTI	rh_caudalanteriorcingulate ↔ right_cerebellum_cortex	416.0x	-2.0×10^{-5}	-1.2×10^{-5}	0.00
DTI	rh_caudalanteriorcingulate ↔ rh_fusiform	393.4x	-1.5×10^{-5}	-1.3×10^{-5}	0.00
DTI	rh_caudalanteriorcingulate ↔ rh_supramarginal	392.8x	-1.3×10^{-5}	-1.0×10^{-5}	0.00
DTI	rh_caudalanteriorcingulate ↔ rh_postcentral	384.6x	-9.0×10^{-6}	-1.1×10^{-5}	0.00
DTI	rh_caudalanteriorcingulate ↔ rh_inferiorparietal	337.0x	-8.0×10^{-6}	-9.0×10^{-6}	0.00
DTI	rh_caudalanteriorcingulate ↔ rh_superiortemporal	317.2x	-1.2×10^{-5}	-6.0×10^{-6}	0.00
DTI	rh_caudalanteriorcingulate ↔ rh_transversetemporal	296.8x	-9.0×10^{-6}	-6.0×10^{-6}	0.00

isolating the true signal from network noise. Table III reports the most salient connections in fMRI and DTI useful for the classification task. The class-specific columns (PD, HC, Prod.) report the mean gradient-weighted feature attributions for each edge. Negative values indicate that a reduction in connectivity strength—representing neurodegenerative structural or functional decay—drives the model toward predicting that specific clinical class

The most important fMRI connections are dominated by the Optic Chiasm and the Anterior Corpus Callosum, with connections to the Frontal Lobe and the Caudate. While the optic chasm itself doesn't generate functional BOLD signals, in the ASEG brain atlas the optic chasm aligns with the anteroventral boundary of the Ventral Diencephalon [10]. The VentralDC contains regions such as the hypothalamus and substantia nigra; two crucial regions in the development of PD. The model learned that breakdowns in communication between the hypothalamus and the frontal cortex, and the decay of functional connections in fronto-striatal loops, are the primary neural biomarkers of PD [11].

The most salient connections in DTI are heavily dominated by the Right Caudal Anterior Cingulate, disconnecting from temporal, parietal, and cerebellar cortices. The latter regions are respectively responsible for processing auditory information and language comprehension, processing sensory information, and regulating motor activity. Their structural isolation from the Right Caudal Anterior Cingulate which has the main task of conflict monitoring [12] (identifying and resolving conflicts from information coming from other brain regions) represents a typical biomarker of PD [13].

V. CONCLUSION

This framework can effectively resolve the complexities of data scarcity and neurodegenerative disease assessment. By successfully imputing missing data and mapping heterogeneous brain graphs into a unified clinical profile, the framework achieves promising predictive accuracy.

Crucially, by offering localized, biologically grounded insights, the framework advances the technical frontier of medical AI. Future work integrating other neuroimaging modalities like SPECT and MRI and extending the framework to downstream tasks like static severity prediction and disease progression forecasting holds immense promise for closing the gap between AI-driven neurodiagnostics and precision medicine.

REFERENCES

[1] W. T and H. M, "The cerebellum in parkinson's disease," vol. 136. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/23404337/

[2] G. Qu, Z. Zhou, V. D. Calhoun, A. Zhang, and Y.-P. Wang, "Integrated brain connectivity analysis with fMRI, DTI, and sMRI powered by interpretable graph neural networks," vol. 103, p. 103570. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1361841525001173

[3] N. L. D. Le, J. Ren, C. Peng, C. Xie, B. Li, and F. Xia, "BrainMAP: Multimodal graph learning for efficient brain disease localization." [Online]. Available: http://arxiv.org/abs/2506.11178

[4] A. Harikumar, M. Misiura, D. Amen, D. Keator, and V. Calhoun, "Fully automated estimation of fMRI-guided SPECT brain networks and their functional network connectivity in schizophrenia patients versus controls: A NeuroMark ICA approach," vol. 5. [Online]. Available: https://apertureneuro.org/article/138726-fully-automated-estimation-of-fmri-guided-spect-brain-networks-and-their-functional-network-connectivity-in-schizophrenia-patients-versus-controls-a

[5] B. P. Rogers, V. L. Morgan, A. T. Newton, and J. C. Gore, "Assessing functional connectivity in the human brain by fmri," *Magnetic resonance imaging*, vol. 25, no. 10, pp. 1347–1357, 2007.

[6] E. Rusak, P. Reizinger, A. Juhos, O. Bringmann, R. S. Zimmermann, and W. Brendel, "Infonce: Identifying the gap between theory and practice," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2025, pp. 4159–4167.

[7] W. Hu, B. Liu, J. Gomes, M. Zitnik, P. Liang, V. Pande, and J. Leskovec, "Strategies for pre-training graph neural networks," in *International Conference on Learning Representations*, 2020.

[8] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?" [Online]. Available: http://arxiv.org/abs/1810.00826

[9] K. Marek, D. Jennings, S. Lasch, A. Siderowf, C. Tanner, T. Simuni, C. Coffey, K. Kieburtz, E. Flagg, S. Chowdhury, W. Poewe, B. Mollenhauer, P.-E. Klinik, T. Sherer, M. Frasier, C. Meunier, A. Rudolph, C. Casaceli, J. Seibyl, S. Mendick, N. Schuff, Y. Zhang, A. Toga, K. Crawford, A. Ansbach, P. De Blasio, M. Piovella, J. Trojanowski, L. Shaw, A. Singleton, K. Hawkins, J. Eberling, D. Brooks, D. Russell, L. Leary, S. Factor, B. Sommerfeld, P. Hogarth, E. Pighetti, K. Williams, D. Standaert, S. Guthrie, R. Hauser, H. Delgado, J. Jankovic, C. Hunter, M. Stern, B. Tran, J. Leverenz, M. Baca, S. Frank, C.-A. Thomas, I. Richard, C. Deeley, L. Rees, F. Sprenger, E. Lang, H. Shill, S. Obradov, H. Fernandez, A. Winters, D. Berg, K. Gauss, D. Galasko, D. Fontaine, Z. Mari, M. Gerstenhaber, D. Brooks, S. Malloy, P. Barone, K. Longo, T. Comery, B. Ravina, I. Grachev, K. Gallagher, M. Collins, K. L. Widnell, S. Ostrowizki, P. Fontoura, T. Ho, J. Luthman, M. V. D. Brug, A. D. Reith, and P. Taylor, "The parkinson progression marker initiative (PPMI)," vol. 95, no. 4, pp. 629–635. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0301008211001651

[10] B. Fischl, D. H. Salat, E. Busa, M. Albert, M. Dieterich, C. Haselgrove, A. Van Der Kouwe, R. Killiany, D. Kennedy, S. Klaveness, A. Montillo, N. Makris, B. Rosen, and A. M. Dale, "Whole Brain Segmentation," *Neuron*, vol. 33, no. 3, pp. 341–355, Jan. 2002.

[11] J. Dukart, F. Sambataro, and A. Bertolino, "Distinct Role of Striatal Functional Connectivity and Dopaminergic Loss in Parkinson's Symptoms," *Frontiers in Aging Neuroscience*, vol. 9, p. 151, May 2017.

[12] S. Ursu, K. A. Clark, H. J. Aizenstein, V. A. Stenger, and C. S. Carter, "Conflict-related activity in the caudal anterior cingulate cortex in the absence of awareness," *Biological Psychology*, vol. 80, no. 3, pp. 279–286, Mar. 2009.

[13] B. A. Vogt, "Cingulate cortex in Parkinson's disease," *Handbook of Clinical Neurology*, vol. 166, pp. 253–266, 2019.