

# Conditional Neural ODE for Longitudinal Parkinson’s Disease Progression Forecasting

Xiaoda Wang\*, Yuji Zhao<sup>‡</sup>, Kaiqiao Han<sup>†</sup>, Xiao Luo<sup>||</sup>, Sanne van Rooij\*, Jennifer Stevens\*,  
Lifang He<sup>§</sup>, Liang Zhan<sup>¶</sup>, Yizhou Sun<sup>†</sup>, Wei Wang<sup>†</sup>, Carl Yang\*

\*Emory University <sup>†</sup>University of California, Los Angeles <sup>‡</sup>Illinois Institute of Technology

<sup>§</sup>Lehigh University <sup>¶</sup>University of Pittsburgh <sup>||</sup>University of Wisconsin–Madison

{xiaoda.wang,sanne.van.rooij,jennifer.stevens,j.carlyang}@emory.edu yzhao104@hawk.iit.edu  
{kqhan,yzsun,weiwang}@cs.ucla.edu xiao.luo@wisc.edu lih319@lehigh.edu liang.zhan@pitt.edu

**Abstract**—Parkinson’s disease (PD) shows heterogeneous, evolving brain-morphometry patterns. Modeling these longitudinal trajectories enables mechanistic insight, treatment development, and individualized ‘digital-twin’ forecasting. However, existing methods usually adopt recurrent neural networks and transformer architectures, which rely on discrete, regularly sampled data while struggling to handle irregular and sparse magnetic resonance imaging (MRI) in PD cohorts. Moreover, these methods have difficulty capturing individual heterogeneity including variations in disease onset, progression rate, and symptom severity, which is a hallmark of PD. To address these challenges, we propose CNODE (Conditional Neural ODE), a novel framework for continuous, individualized PD progression forecasting. The core of CNODE is to model morphological brain changes as continuous temporal processes using a neural ODE model. In addition, we jointly learn patient-specific initial time and progress speed to align individual trajectories into a shared progression trajectory. We validate CNODE on the Parkinson’s Progression Markers Initiative (PPMI) dataset. Experimental results show that our method outperforms state-of-the-art baselines in forecasting longitudinal PD progression.

**Index Terms**—Neural ODE, Parkinson’s disease progression, Time series forecasting, Contrastive learning, MRI.

## I. INTRODUCTION

PD is a common neurodegenerative disorder with heterogeneous motor and non-motor manifestations and substantial societal burden [1], [2]. Clinically, PD presents with a wide spectrum of motor symptoms and non-motor symptoms [6], [15]. The progression of PD along with these symptoms varies across individuals, which reflects heterogeneous pathophysiological changes over years [17]–[19]. And magnetic resonance imaging (MRI) has become a valuable tool to study morphological alterations associated with PD [10], [20]. Recent advances [9] in structural and functional MRI analyses have revealed PD’s neuroanatomical correlates, including changes in deep brain structures, and shifts in functional connectivity.

However, capturing the full PD progression is challenging. In particular, we expect effective analytical methods that can model continuous, and highly individualized trajectories of brain morphology, which remain difficult due to the nature of PD. First, patients undergo MRI at varying intervals in real-world clinical settings. Traditional deep learning methods such as recurrent neural networks (RNNs) and transformer-based architectures typically assume regularly sampled data, making

them poorly suited for handling inherently continuous data sampled with varying intervals or missing values. Second, the disease onset age, progression rate, and symptom severity vary greatly among patients, which indicates distinct clinical trajectories even in relatively homogeneous cohorts [12]. When studying longitudinal progression, it is highly anticipated to capture the intrinsic heterogeneity of PD progression.

To address these challenges, we propose CNODE (Conditional Neural ODE), a novel framework designed to forecast longitudinal PD progression from irregular and sparsely sampled MRI scans. Specifically, CNODE models disease evolution as a continuous trajectory in a shared latent space, while jointly learning patient-specific initial time and progression speed as conditioning variables. Then CNODE captures individualized disease dynamics and effectively aligns all patients onto a unified progression trajectory, enabling robust prediction and analysis of heterogeneous disease courses.

In summary, our main contributions are as follows:

- We are the first to recognize the importance of modeling the continuous progression for PD and propose CNODE to model continuous PD progression, effectively addressing the challenges of irregular sampled MRI data.
- We harness subcortical segmentations from T1-weighted MRI to extract vertex-specific shape features. Instead of modeling the subject-wise whole brain shape directly, CNODE extracts the vertex-wise medial thickness.
- CNODE aligns all patient into a shared progression trajectory by jointly learning patient-specific initial time and progression speed. CNODE also leverages contrastive learning for trajectory alignment.
- We validate CNODE on the PPMI dataset, showing that our method outperforms state-of-the-art baselines and provides insights into personalized disease trajectories forecasting.

## II. METHODOLOGY

In this section, we present CNODE, a Conditional Neural ODE-based framework to model the longitudinal PD progression as shown in Figure 1. First, we extract vertex-wise *medial thickness* from MRI scans, with subcortical structure volumes and metadata as condition features. Second, to address sparse visit times, we construct a shared progression trajectory. We

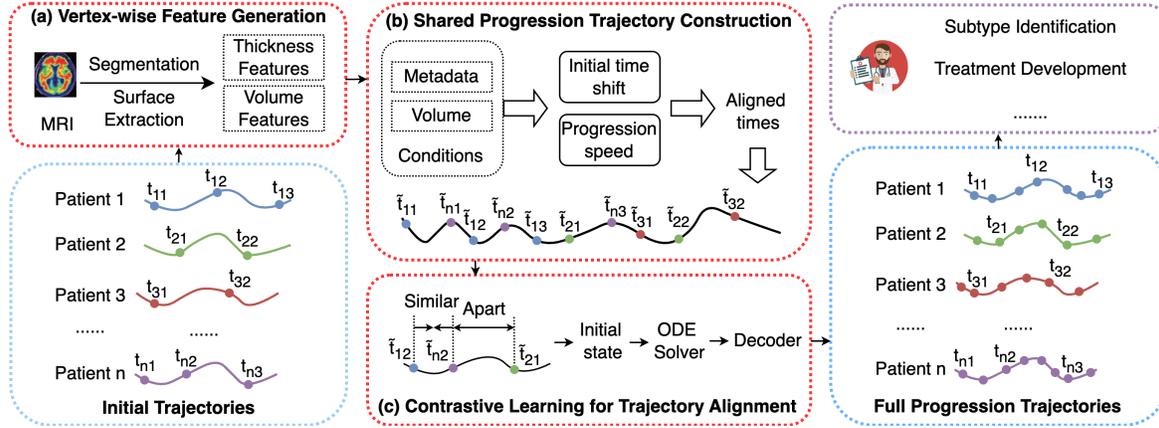


Fig. 1: Overview of CNODE. (a) Firstly, we extract vertex-wise *medial thickness* from MRI, with subcortical structure volumes. (b) Secondly, we construct a shared progression trajectory. (c) Then we employ contrastive learning to align time distributions. Finally, we model the continuous evolution with Neural ODE.

also employ contrastive learning, leveraging shape feature similarity to enforce aligned time distributions. Then we model the continuous evolution of shape features with Neural ODE [13]. **Problem Definition.** We focus on the problem of continuous and longitudinal PD progression forecasting. Let us consider a set of  $N$  subjects (patients) denoted by  $S = \{s_i\}_{i=1}^N$ . Each subject  $s_i$  has  $K_i$  observed timestamps  $\mathcal{T}_i = \{t_{i,1}, t_{i,2}, \dots, t_{i,K_i}\}$ , where these timestamps may be non-uniformly spaced and can take arbitrary continuous values. At each time  $t_{i,k}$ , we extract *shape features*  $X_{i,k} \in \mathbb{R}^D$ , where  $D$  is the dimensionality of the feature vector. Additionally, each subject  $s_i$  has associated *volume features*  $C_i$  and *metadata*  $M_i$ . Given a limited number of observed shape features  $\{X_{i,0:n_i}\}$  for each subject  $s_i$ , along with the subject-specific volume features  $C_i$  and metadata  $M_i$ , our goal is to forecast the future shape features  $\{X_{i,n_i+1:T_i}\}$  at arbitrary future times. Formally, we aim to learn a model  $\hat{X}_{i,n_i+1:T_i} = f(\{X_{i,0:n_i}\}, C_i, M_i; \theta)$ , where  $\theta$  represents the model parameters.

### A. Vertex-wise Feature Generation

Figure 1 illustrates our subcortical mesh-based feature generation process. We extract morphometrical features with respect to subcortical and cortical brain structures. We focus on *medial thickness* as the primary shape feature. Cortical medial thickness is extracted at each vertex, and the average thickness is computed across all vertices within each of the 34 predefined bilateral cortical per hemisphere, as defined by the Desikan-Killiany atlas. The resulting vertex-wise thickness values at time  $t_{i,k}$  form the shape feature vector  $X_{i,k}$  for subject  $s_i$ . We also extract *subcortical structure volumes*. Given that some subjects have only one or two visits, the volumes features  $C_i$  and metadata  $M_i$  (age, gender) from the first visit serve as the condition features.

### B. Shared Progression Trajectory Construction

Longitudinal visit times are sparse and irregular, which impedes learning a reliable disease timeline. We leverage

subcortical volumes and metadata—informative for PD staging [8], [21]—to *align* all subjects onto a shared progression axis by estimating (i) an onset shift and (ii) an individual speed. Two lightweight regressors  $f_{\theta_\tau}$  and  $f_{\theta_\gamma}$  map patient metadata  $M_i$  and condition features  $C_i$  to these alignment parameters:

$$\tau_i = f_{\theta_\tau}(M_i, C_i), \quad \gamma_i = f_{\theta_\gamma}(M_i, C_i). \quad (1)$$

Given a raw visit time  $t_{i,k}$  and the subject’s first visit  $t_{i,1}$ , the aligned time is

$$\tilde{t}_{i,k} = \tau_i + \gamma_i (t_{i,k} - t_{i,1}), \quad (2)$$

which standardizes trajectories by correcting onset heterogeneity and scaling progression rates.

### C. Contrastive Learning for Shared Trajectory Alignment

Aligned times lack ground truth. To regularize them, we impose a metadata-aware contrastive objective so that visits with *similar shape features* are temporally close, and dissimilar ones are far. We declare a pair  $(i, k)$  and  $(j, \ell)$  as *positive* if the cosine similarity between their shape features  $X_{i,k}$  and  $X_{j,\ell}$  exceeds a threshold  $\delta$ ; otherwise it is *negative*. We weight each pair by the similarity of their conditions:

$$w_{ij} = \text{sim}((M_i, C_i), (M_j, C_j)). \quad (3)$$

Let  $s(\tilde{t}_i, \tilde{t}_j)$  denote cosine similarity between aligned times. Using temperature  $\tau$ , the metadata-weighted InfoNCE loss is

$$\mathcal{L}_{\text{contrast}} = - \sum_i \sum_{j: \text{positive to } i} w_{ij} \log \frac{\exp(s(\tilde{t}_i, \tilde{t}_j)/\tau)}{\sum_k \exp(s(\tilde{t}_i, \tilde{t}_k)/\tau)}. \quad (4)$$

This encourages anchors to be close to positive visits while contrasting against all others, with stronger influence from clinically similar patients via  $w_{ij}$ . The result is a shared progression timeline that reflects both morphology and patient condition under sparse, irregular sampling.

TABLE I: **Comparison of PD progression forecasting performance.** All metrics are presented as averages  $\pm$  standard deviations under 5-fold cross-validation. The best performances are in bold.

Model	MSE $\downarrow$	RMSE $\downarrow$	$R^2$ $\uparrow$
GRU	0.0281 $\pm$ 0.0012	0.1675 $\pm$ 0.0036	0.8112 $\pm$ 0.0082
RNN	0.0305 $\pm$ 0.0027 $\uparrow$ 0.0024	0.1745 $\pm$ 0.0076 $\uparrow$ 0.0070	0.7947 $\pm$ 0.0179 $\downarrow$ 0.0165
LSTM	0.0292 $\pm$ 0.0019 $\uparrow$ 0.0011	0.1708 $\pm$ 0.0056 $\uparrow$ 0.0033	0.8037 $\pm$ 0.0130 $\downarrow$ 0.0075
Transformer	0.0291 $\pm$ 0.0010 $\uparrow$ 0.0010	0.1705 $\pm$ 0.0030 $\uparrow$ 0.0030	0.8038 $\pm$ 0.0069 $\downarrow$ 0.0074
Neural ODE	0.0283 $\pm$ 0.0019 $\uparrow$ 0.0002	0.1632 $\pm$ 0.0043 $\downarrow$ 0.0043	0.7907 $\pm$ 0.0147 $\downarrow$ 0.0205
LLMTime	0.0282 $\pm$ 0.0105 $\uparrow$ 0.0001	0.1652 $\pm$ 0.0226 $\downarrow$ 0.0023	0.8041 $\pm$ 0.0739 $\downarrow$ 0.0071
CNODE (Ours)	<b>0.0258 <math>\pm</math> 0.0001</b> $\downarrow$ 0.0023	<b>0.1606 <math>\pm</math> 0.0003</b> $\downarrow$ 0.0069	<b>0.8260 <math>\pm</math> 0.0006</b> $\uparrow$ 0.0148

#### D. Neural ODE for Longitudinal Progression Modeling

Given the aligned visit times  $\tilde{t}_{i,k}$ , we model subject-specific continuous trajectories in a latent space and decode them back to shape features.

1) *Initial State Encoder*: For subject  $i$ , the encoder summarizes the observed trajectory into a Gaussian posterior over the initial latent state:

$$q_\phi(z_{i,0} | X_{i,0:n_i}) = \mathcal{N}(\mu_i, \text{diag}(\sigma_i^2)) \quad (5)$$

$$(\mu_i, \log \sigma_i) = \text{Enc}_\phi(X_{i,0:n_i}). \quad (6)$$

2) *ODE Generative Model*: Latent dynamics evolve continuously along the aligned timeline  $\tilde{\mathbf{t}}_i = \{\tilde{t}_{i,1}, \dots, \tilde{t}_{i,n_i}\}$ :

$$z_i(\tilde{\mathbf{t}}_i) = \text{ODEsolve}(g_\theta, z_{i,0}, \tilde{\mathbf{t}}_i), \quad (7)$$

where  $g_\theta$  is a neural ODE field parameterizing disease progression.

3) *Decoder*: At each aligned time, the decoder reconstructs shape features:

$$\hat{X}_{i,k} = D_\theta(z_i(\tilde{t}_{i,k})) \quad (8)$$

$$p(X_{i,k} | z_i(\tilde{t}_{i,k})) = \mathcal{N}(\hat{X}_{i,k}, \sigma_x^2 I). \quad (9)$$

This yields flexible, continuous-time forecasts under sparse and irregular sampling.

#### E. Training

We jointly optimize the encoder, ODE field, and decoder end-to-end with reconstruction, regularization, and alignment terms. Over all subjects and observed visits:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{\sum_i n_i} \sum_i \sum_{k=1}^{n_i} \|X_{i,k} - \hat{X}_{i,k}\|_2^2, \quad (10)$$

$$\mathcal{L}_{\text{KL}} = \sum_i D_{\text{KL}}(\mathcal{N}(\mu_i, \text{diag}(\sigma_i^2)) \| \mathcal{N}(0, I)), \quad (11)$$

The final objective is

$$\mathcal{L} = \mathcal{L}_{\text{MSE}} + \lambda_{\text{KL}} \mathcal{L}_{\text{KL}} + \lambda_{\text{C}} \mathcal{L}_{\text{contrast}}, \quad (12)$$

with  $\lambda_{\text{KL}}$  and  $\lambda_{\text{C}}$  balancing regularization and alignment.

### III. EXPERIMENT

In this section, we comprehensively evaluate the effectiveness of the proposed CNODE framework through extensive experiments on longitudinal Parkinson’s disease (PD) data.

#### A. Dataset Description and Implementation Details

1) *Dataset Description*: We evaluate CNODE using the PPMI dataset [14], a comprehensive, multi-center study aimed at identifying progression markers for Parkinson’s disease. For this study, we utilize T1-weighted MRI, which provides high-resolution anatomical details of the brain. The vertex-wise features were extracted from 68 sub-cortical brain regions using FreeSurfer. The final cohort includes 161 PD subjects, with 50 individuals having three visits and another 111 having two visits. The average visit interval is 1.11 years, with a maximum of 2.27 years and a minimum of 0.61 years.

2) *Evaluation Metrics*: The effectiveness of each model was evaluated using Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and the Coefficient of Determination ( $R^2$ ). MSE quantifies prediction error, RMSE provides magnitude, and  $R^2$  measures the variance explained. Lower MSE and RMSE indicate better accuracy, while a higher  $R^2$  suggests improved trend capture.

#### B. Overall PD Progression Forecasting Performance

We evaluated our model’s performance in forecasting PD progression against several baseline models:

- **Recurrent Neural Network (RNN)** [5]: A classical sequential model capturing temporal dependencies.
- **Long Short-Term Memory (LSTM)** [11] and **Gated Recurrent Unit (GRU)** [4]: Popular RNN variants designed to address long-term dependency and vanishing gradient issues.
- **Transformer** [16]: A self-attention-based model that can flexibly model long-range dependencies in temporal sequences.
- **Neural Ordinary Differential Equation (Neural ODE)** [3]: A continuous-time latent dynamics model for irregularly sampled time series.
- **LLMTime** [7]: Uses large language models to perform time series forecasting in a zero-shot manner.

For baselines designed for regularly spaced intervals, we set the interval to one year. Table I summarizes the forecast results for all models. As shown, CNODE consistently achieves superior performance across all evaluation metrics, including the lowest Mean Squared Error (MSE), lowest Root Mean Squared Error (RMSE), and highest coefficient of determination ( $R^2$ ). These results highlight the effectiveness of our approach in capturing the continuous and heterogeneous

progression of PD, even under irregular and sparse sampling conditions.

### C. Hyperparameter Analysis

We performed a hyperparameter analysis to examine how the balance between the KL divergence loss ( $\lambda_1$ ) and the contrastive loss ( $\lambda_2$ ) affects model performance. As shown in Figure 2, we varied the ratio  $\lambda_1/\lambda_2$  across a broad range. The results demonstrate that CNODE achieves the best forecasting accuracy—reflected by lower MSE/RMSE and higher  $R^2$ —when  $\lambda_1/\lambda_2$  is set to 1:1. These findings highlight the importance of properly weighting the KL and contrastive objectives for learning effective latent representations and robust trajectory alignment.

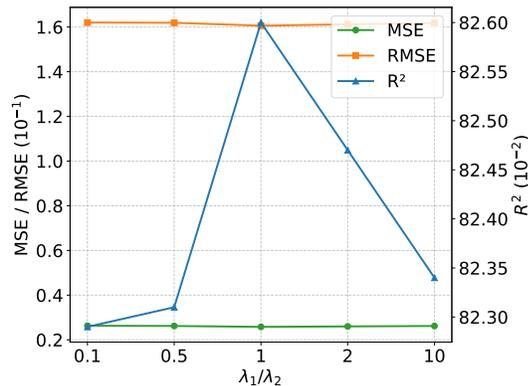


Fig. 2: Effect of  $\lambda_1/\lambda_2$  on CNODE performance.

## IV. CONCLUSION

In this work, we introduced CNODE, a Conditional Neural ODE framework that models continuous Parkinson’s disease progression. By harnessing subcortical shape features and aligning patients within a shared progression trajectory, CNODE captures individualized disease dynamics while addressing common challenges such as missing timepoints and heterogeneous progression rates. Experimental results demonstrate the superiority of our proposed CNODE over state-of-the-art methods, highlighting its potential for personalized disease monitoring and therapeutic decision support.

## V. ACKNOWLEDGEMENT

This research was partially supported by the U.S. National Science Foundation under Award Numbers 2319449, 2312502, 2442172, 2045848, 2215789, 2319450, 2319451, 2106859, 2200274, 2312501, 2211557, 2119643, 2303037 and 2531008; by the U.S. National Institutes of Health under Award Numbers U54HG012517, U24DK097771, U54OD036472, and OT2OD038003; and by the US National Institute of Diabetes and Digestive and Kidney Diseases of the US National Institutes of Health under Award Number K25DK135913. Additional support was provided by SRC JUMP 2.0 Center, Amazon, NEC, Optum AI, Amazon Research Awards, and Snapchat Gifts.

## REFERENCES

- [1] Aarsland, D., Batzu, L., Halliday, G.M., Geurtsen, G.J., Ballard, C., Ray Chaudhuri, K., Weintraub, D.: Parkinson disease-associated cognitive impairment. *Nature reviews Disease primers* 7(1), 47 (2021)
- [2] Armstrong, M.J., Okun, M.S.: Diagnosis and treatment of parkinson disease: a review. *Jama* 323(6), 548–560 (2020)
- [3] Chen, R.T., Rubanova, Y., Bettencourt, J., Duvenaud, D.K.: Neural ordinary differential equations. *Advances in neural information processing systems* 31 (2018)
- [4] Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y.: Learning phrase representations using rnn encoder-decoder for statistical machine translation. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. pp. 1724–1734 (2014)
- [5] Elman, J.L.: Finding structure in time. *Cognitive science* 14(2), 179–211 (1990)
- [6] Gratwicke, J., Jahanshahi, M., Foltyniec, T.: Parkinson’s disease dementia: a neural networks perspective. *Brain* 138(6), 1454–1476 (2015)
- [7] Gruver, N., Finzi, M., Qiu, S., Wilson, A.G.: Large language models are zero-shot time series forecasters. *Advances in Neural Information Processing Systems* 36, 19622–19635 (2023)
- [8] Gu, W., He, R., Su, H., Ren, Z., Zhang, L., Yuan, H., Zhang, M., Ma, S.: Changes in the shape and volume of subcortical structures in patients with end-stage renal disease. *Frontiers in Human Neuroscience* 15, 778807 (2021)
- [9] Han, K.M., Ham, B.J.: How inflammation affects the brain in depression: a review of functional and structural mri studies. *Journal of clinical neurology (Seoul, Korea)* 17(4), 503 (2021)
- [10] Herold, F., Wiegel, P., Scholkmann, F., Thiers, A., Hamacher, D., Schega, L.: Functional near-infrared spectroscopy in movement science: a systematic review on cortical activity in postural and walking tasks. *Neurophotonics* 4(4), 041403–041403 (2017)
- [11] Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural computation* 9(8), 1735–1780 (1997)
- [12] Lian, J., Luo, X., Shan, C., Han, D., Zhang, C., Vardhanabhuti, V., Li, D., Qiu, L.: Personalized progression modelling and prediction in parkinson’s disease with a novel multi-modal graph approach. *npj Parkinson’s Disease* 10(1), 229 (2024)
- [13] Liu, Z., Wang, X., Wang, B., Huang, Z., Yang, C., Jin, W.: Graph odes and beyond: A comprehensive survey on integrating differential equations with graph neural networks. In: *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*. pp. 6118–6128 (2025)
- [14] Marek, K., Jennings, D., Lasch, S., Siderowf, A., Tanner, C., Simuni, T., Coffey, C., Kiebertz, K., Flagg, E., Chowdhury, S., et al.: The parkinson progression marker initiative (ppmi). *Progress in neurobiology* 95(4), 629–635 (2011)
- [15] Sousani, M., Seydnejad, S.R., Ghahramani, M.: Insights from a model based study on optimizing non invasive brain electrical stimulation for parkinson’s disease. *Scientific Reports* 14(1), 2447 (2024)
- [16] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: *Advances in neural information processing systems*. pp. 5998–6008 (2017)
- [17] Wielinski, C.L., Erickson-Davis, C., Wichmann, R., Walde-Douglas, M., Parashos, S.A.: Falls and injuries resulting from falls among patients with parkinson’s disease and other parkinsonian syndromes. *Movement Disorders* 20(4), 410–415 (2005)
- [18] Xie, Y., Cui, H., Zhang, Z., Lu, J., Shu, K., Nahab, F., Hu, X., Yang, C.: Kerap: A knowledge-enhanced reasoning approach for accurate zero-shot diagnosis prediction using multi-agent llms. In: *American Medical Informatics Association 2025 Annual Symposium* (2025)
- [19] Xie, Y., Han, X., Xu, R., Hu, X., Lu, J., Yang, C.: Hypkg: Hypergraph-based knowledge graph contextualization for precision healthcare. In: *24th International Semantic Web Conference* (2025)
- [20] Zhao, Y., Gutman, B.A.: Morph-transformer: A deep spatiotemporal framework for shape analysis in longitudinal neuroimaging. In: *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*. pp. 1–5. IEEE (2024)
- [21] Zhao, Y., Laansma, M.A., van Heese, E.M., Owens-Walton, C., Parkes, L.M., Debove, I., Rummel, C., Wiest, R., Cendes, F., Guimaraes, R., et al.: Learning interpretable regularized ordinal models from 3d mesh data for neurodegenerative disease staging. *Machine Learning in Clinical Neuroimaging, MLCN 2022* 13596, 115–124 (2022)