



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

ENC-ODE: Event-level Neurodegenerative Modeling in Continuous Time with Neural ODEs

Yujee Song^{1*}, Seunghun Baek^{1*}, Guorong Wu², and Won Hwa Kim¹

¹ Pohang University of Science and Technology, Pohang, South Korea

{yujees, habaek4, wonhwa}@postech.ac.kr

² University of North Carolina at Chapel Hill, Chapel Hill, USA

guorong_wu@med.unc.edu

Abstract. Accurately predicting the temporal evolution of clinical biomarkers is crucial for the early diagnosis and management of neurodegenerative diseases such as Alzheimer’s disease. However, this relies on longitudinal data to capture biomarker changes over time, which is often sparse and irregular due to the high cost, labor-intensive nature, and patient burden. To address these challenges, we propose ENC-ODE, an Event-level Neurodegenerative modeling in Continuous time with neural Ordinary Differential Equations. ENC-ODE predicts future biomarker evolution by modeling clinical events through diagnosis-conditioned continuous dynamics. A target-conditioned attention mechanism weights and aggregates event-level predictions for the target time and modality without history compression. Extensive experiments on Alzheimer’s Disease Neuroimaging Initiative (ADNI) dataset demonstrate that ENC-ODE outperforms representative sequence models while offering a scalable and neuroscientifically grounded solution for clinical support. The code is available at <https://github.com/JardinDelSol/enc-ode>.

Keywords: Alzheimer’s Disease · Time Series Analysis · Neuroimaging

1 Introduction

Accurately modeling temporal dynamics of human brain is essential for improving patient management and early diagnosis of neurodegenerative diseases [12, 2]. For example, predicting how biomarkers such as Standardized Uptake Value Ratio (SUVR) of β -amyloid, metabolism, and Tau protein evolve over time helps clinicians to identify early symptoms of Alzheimer’s disease (AD) and tailored treatment planning [7, 10, 26]. However, it is difficult to obtain a reliable prediction, as longitudinal neurobiological biomarkers are inherently sparse due to the high cost and patient burden from imaging scans [3, 4]. Thus, developing novel methods to predict future biomarker states from sparsely observed data is critical, as it provides a cost-effective alternative to radiation-based examinations.

Recent efforts explore diverse sequence modeling for neurodegenerative disease forecasting, which generally follow two primary paradigms categorized into

* Y. Song and S. Baek contributed equally to this paper.

discrete-time and continuous-time architectures. Discrete-time architectures including RNN [24], Transformer [27], and Mamba [13] are designed for regular intervals. As a result, they struggle to capture the continuous dynamics of disease progression under irregular time gaps. Alternatively, continuous-time models such as Neural Ordinary Differential Equations (ODEs) [8], ODE-RNN [23], Neural Flow [5], and Continuous Recurrent Unit (CRU) [25] handle temporal irregularity. Despite their different approaches to time, both paradigms share a structural limitation that they predominantly compress entire multimodal histories into a single summarized latent representation. Such compression fails to capture explicit dependencies between specific past *events* (e.g., individual modality-specific biomarker observations acquired at irregular time points, along with the corresponding diagnostic stage) and the target time or modality. This structural bottleneck hinders learning how specific target predictions should selectively draw information from highly heterogeneous past observations, which ultimately limits the precise characterization of long-term disease trajectories.

To overcome these limitations, we propose **Event-level Neurodegenerative Continuous-time modeling with neural ODEs (ENC-ODE)**. ENC-ODE propagates each past observation through diagnosis-conditioned ODE dynamics and aggregates the resulting predictions using attention conditioned on the target time point and modality. This design explicitly models biomarker evolution under varying time intervals while avoiding the loss of event level information caused by compressing multimodal histories into a single latent state. By computing attention scores on unobserved target event and dynamically propagated predictions, our model prioritizes meaningful information over restrictive history compression. This design effectively handles irregularly sampled data and provides precise characterization of diagnosis-specific biomarker trajectories.

Our **contributions** are summarized as follows: **i) *Diagnosis-conditioned event-level continuous trajectory modeling***: We propose diagnosis-driven dynamics for each individual observed event to capture stage-specific progression rates across irregular time intervals, enabling the precise characterization of continuous biomarker evolution. **ii) *Target-conditioned multimodal aggregation***: We introduce an attention module that selectively aggregates past observations based on their direct relevance to the target time and modality, bypassing the information loss inherent in previous unified latent bottlenecks. **iii) *Extensive validation***: We demonstrate the effectiveness of ENC-ODE via extensive analyses, validating its generation of plausible biomarker trajectories (i.e., changes of SUVR measures) on Alzheimer’s Disease Neuroimaging Initiative (ADNI) dataset.

2 Methods

Our approach aims to model longitudinal recordings on multimodal brain imaging and predict the temporal evolution of the brain with respect to disease stages. We denote the longitudinal data as $\mathcal{H}_{t_n} := \{e_{t_i}\}_{i=1}^n$, where $t_1 < t_2 < \dots < t_n$ are ordered time points and e_i represents an event (e.g., image acquisition) at time t_i . Each event $e_{t_i} := \{m_{t_i}, f_{t_i}, \delta_{t_i}, t_i\}$ consists of the modality type m_{t_i}




Fig. 1: **Overview of ENC-ODE.** The input consists of multimodal brain regional measures $f(t_i)$ collected at irregular time points t_i . The $f(t_i)$ is first encoded into a hidden state h where Neural ODEs model the continuous evolution of hidden states $\hat{h}$. The learned trajectories capture biomarker dynamics, enabling the prediction of future states. The final prediction $\hat{f}(t_p)$ aggregates event-level predictions via weighted summation, reflecting the relative contribution of each prior event to the target prediction.

(e.g., SUVR of β -amyloid), the measurement from N regions of interests (ROIs) $f_{t_i} \in \mathbb{R}^N$, the diagnosed AD stages δ_{t_i} , and the observed time point t_i .

Similar to RNN and Transformer-based models, we first encode e_{t_i} into a hidden state $h(t|e_{t_i}) \in \mathbb{R}^{dim_h}$, but we disentangle event-driven changes over time using ODEs as shown in Fig. 1. Each observation captures different aspects of the subject's condition, and ODEs are used to model how the ROIs evolve based on the patient's state and diagnosis. The predictions for the follow-up measurement, $\hat{f}(t_p)$, are made by solving the initial value problem [21] (IVP) and aggregating the resulting hidden states that incorporate prior events $\{e_{t_i}\}_{i=1}^{p-1}$.

2.1 Hidden State Propagation

As we consider observations made from different modalities, we first map our observations into latent representations. Specifically, the event e_{t_i} is encoded into hidden state, $h(t_i|e_{t_i})$, using a neural network $g_{enc}(\cdot)$:

$$h(t_i|e_{t_i}) = g_{enc}(e_{t_i}) = g_{enc}([f_{t_i} \parallel \xi_m(m_{t_i}) \parallel \xi_\delta(\delta_{t_i})]), \quad (1)$$

where $\xi_m(\cdot) : \mathbb{R} \rightarrow \mathbb{R}^{dim_m}$ and $\xi_\delta(\cdot) : \mathbb{R} \rightarrow \mathbb{R}^{dim_\delta}$ are learnable mappings that transform modality and diagnosis into dim_m and dim_δ dimensional vectors.

Brain measurements exhibit natural temporal fluctuations driven by aging and disease progression. Therefore, the hidden state representing the patient's condition should also evolve over time. To address this, our approach captures these evolutions by modeling the hidden state dynamics using an ODE:

$$dh(t|e_{t_i}) = \gamma(h(t|e_{t_i}), \xi_\delta(\delta_{t_i}); \theta) dt \quad (2)$$

$$h(t_p|e_{t_i}) = h(t_i|e_{t_i}) + \int_{t_i}^{t_p} \gamma(h(s|e_{t_i}), \xi_\delta(\delta_{t_i}); \theta) ds \quad (3)$$

where $\gamma(\cdot)$ is a Neural ODE module parameterized by a neural network θ . The function $\gamma(\cdot)$ models the evolution of the patient’s state over time while incorporating diagnosis information δ , enabling the model to learn diagnosis-conditioned dynamics, as individuals in different disease stage may exhibit heterogeneous dynamics. By solving the IVP of this ODE, we obtain the hidden state trajectory, which captures how the target hidden state $h(\cdot)$ evolves over time. Neural ODEs effectively handle sparse and irregular time points [8,23], making them suitable for modeling hidden state dynamics between irregular observations.

Notably, $\gamma(\cdot)$ modeling the rate of change takes only its current state $h(s|e_{t_i})$ and the diagnosis $\xi_\delta(\delta_{t_i})$ as input. This implies that the entire process is conditioned solely on the initial event e_{t_i} independent from other observations. The independent propagation of hidden states makes the model flexible as it allows each trajectory to evolve without being constrained by previous or future observations. This differentiates our method from other approaches that focus on finding a representation summarizing past observation based on inter-event interactions, allowing us to assess the impact of each event on the predicted measurements in the future t_p . Further details are discussed in Sec. 3.2.

Efficient Computation. Despite all the benefits of ODEs, their autoregressive nature may lead to high computational costs. To mitigate this issue, we model the rate of change as a multidimensional system of ODEs:

$$\frac{d}{dt} \mathbf{h}(t_p) = \frac{d}{dt} \begin{bmatrix} h(t_p|e_{t_1}) \\ \vdots \\ h(t_p|e_{t_{p-1}}) \end{bmatrix} = \begin{bmatrix} \gamma(h(t_p|e_{t_1}), \xi_\delta(\delta_{t_1})) \\ \vdots \\ \gamma(h(t_p|e_{t_{p-1}}), \xi_\delta(\delta_{t_{p-1}})) \end{bmatrix}. \quad (4)$$

Eq. (4) computes rate of change of $h(t_p|e_{t_i})$ at time t_p in parallel per e_{t_i} , reducing the complexity of handling a sequence with n observations from $O(n^2)$ to $O(n)$.

2.2 Estimation on Multi-modal Biomarkers from Prior Observations

Given a propagated hidden state $h(t_p|e_{t_i})$, we can obtain a prediction on the measurements $f(t_p)$ at time t_p given the event e_{t_i} as follows:

$$\hat{f}(t_p|e_{t_i}) = g_{dec}(h(t_p|e_{t_i})), \quad (5)$$

where $g_{dec} : \mathbb{R}^{dim_h} \rightarrow \mathbb{R}^N$ is the decoder network that maps the hidden state into the measurements of ROIs. Moreover, to ensure the most reliable estimation, all predictions from prior observations are aggregated. In our work, the final prediction $\hat{f}(t_p)$ is obtained by combining predictions from earlier observations $\{\hat{f}(t_p|e_{t_i})\}_{i=1}^{p-1}$, effectively capturing information from the entire sequence. To explicitly map the specific contribution of each past observation, we aggregate all prior event-level trajectories through a weighted linear summation as

$$\hat{f}(t_p|\mathcal{H}_{t_{p-1}}) = \sum_{t_i < t_p} w_{i,p} \cdot \hat{f}(t_p|e_{t_i}), \quad (6)$$

where $w_{i,p}$ represents the contribution of each past observation to ensure that $\hat{f}(t_p|\mathcal{H}_{t_{p-1}})$ effectively blends information across the entire sequence.

We leverage an attention mechanism to capture the complex relationships between the target prediction time and modality, allowing us to identify the most influential factors. The value of each $w_{:,p}$ is calculated as follows:

$$w_{:,p} = \text{softmax} \left(\frac{Q([\mathbf{h}(t_p)|\xi_t(t_{1:p-1})])K([\xi_t(t_p)|\xi_m(m_p)])^\top}{\sqrt{\dim_h}} \right), \quad (7)$$

where $\mathbf{h}(\cdot)$ is a multidimensional hidden state from Eq.(4), and $Q(\cdot)$ and $K(\cdot)$ are single-layer linear models that generate the query and key matrices, respectively. Instead of employing a traditional self-attention mechanism, which is commonly used for next-event prediction, we compute attention between our predicted values and the target time and modality. This approach enables the model to emphasize predictions $\hat{f}(t_p|e_{t_i})$ that are most relevant for estimating e_{t_p} , making it particularly effective for handling irregularly sampled time points [28]. The resulting attention weights provide insights into inter-modality relationships, enhancing cross-modal synergies, as further discussed in Section 3.2.

3 Experiment and Results

3.1 Experiment Setup

Dataset. We use ADNI data [20] to validate our framework. Region-specific Standardized Uptake Value Ratio (SUVR) [26] of β -amyloid protein (AMY), metabolism (FDG), and Tau protein (TAU) from positron emission tomography (PET) scans are obtained at $N=160$ brain regions based on Destrieux atlas [11]. The dataset includes 5676 records from 1176 subjects, detailed in Tab. 1. Modality m_t consists of $\{AMY, FDG, TAU\}$, and diagnosis δ_t represents the ordinal AD progression as Cognitive Normal (CN), Subjective Memory Complaint (SMC), Early Mild Cognitive Impairment (EMCI), Late MCI (LMCI) and AD. Events are grouped by subjects and ordered chronologically.

Implementation Details. We use a fully connected layer (FCL) for g_{enc} and g_{dec} . The ODE solver [8] uses the Euler method [6] with a step size of 4, and its structure consists of two FCLs with activation normalization [17]. The model is trained using root mean squared error (RMSE). We train the model for 1000

Table 1: Sample-size of the ADNI dataset categorized by disease labels based on the first record.

Biomarker	Category	CN	SMC	EMCI	LMCI	AD	Total
FDG	# of subjects	255	4	189	277	115	840
	# of records	782	4	455	1016	607	2864
TAU	# of subjects	60	47	49	31	21	208
	# of records	140	112	122	76	56	506
AMY	# of subjects	220	89	237	158	56	760
	# of records	672	240	741	379	274	2306

Table 2: Averaged length of sequence in both scenarios.

Biomarker	Unimodal	Multimodal
FDG	3.42	5.98
TAU	2.43	7.31
AMY	3.03	6.06

Table 3: Comparison of Prediction Scores. Results are reported as the mean with the standard deviation over five independent runs. **Bold face** and underline denote the best performance within each setting and overall, respectively.

	Target	FDG			TAU			AMY		
	Metric	RMSE ↓	MAE ↓	R^2 ↑	RMSE ↓	MAE ↓	R^2 ↑	RMSE ↓	MAE ↓	R^2 ↑
Unimodal	RNN	0.1034 $\pm$ 0.0017	0.078 $\pm$ 0.0012	0.7739 $\pm$ 0.0072	0.2428 $\pm$ 0.0054	0.1338 $\pm$ 0.0031	0.7141 $\pm$ 0.0127	0.2059 $\pm$ 0.0007	0.1561 $\pm$ 0.0005	0.7276 $\pm$ 0.0020
	Transformer	0.1039 $\pm$ 0.0014	0.0779 $\pm$ 0.0008	0.7769 $\pm$ 0.0060	0.2225 $\pm$ 0.0041	0.1258 $\pm$ 0.0014	0.7599 $\pm$ 0.0088	0.1998 $\pm$ 0.0013	0.1506 $\pm$ 0.0010	0.7433 $\pm$ 0.0034
	Neural ODE	0.1068 $\pm$ 0.0013	0.0706 $\pm$ 0.0006	0.7642 $\pm$ 0.0058	0.1681 $\pm$ 0.0030	0.1105 $\pm$ 0.0020	0.8630 $\pm$ 0.0049	0.1897 $\pm$ 0.0010	0.1433 $\pm$ 0.0009	0.7687 $\pm$ 0.0025
	ODE-RNN	0.1027 $\pm$ 0.0010	0.0768 $\pm$ 0.0007	0.7827 $\pm$ 0.0041	0.2211 $\pm$ 0.0049	0.1287 $\pm$ 0.0022	0.7631 $\pm$ 0.0107	0.1901 $\pm$ 0.0023	0.1427 $\pm$ 0.0019	0.7677 $\pm$ 0.0057
	Neural Flows	0.1089 $\pm$ 0.0003	0.0703 $\pm$ 0.0002	0.7552 $\pm$ 0.0016	0.1613 $\pm$ 0.0036	0.1076 $\pm$ 0.0020	0.8738 $\pm$ 0.0056	0.1887 $\pm$ 0.0012	0.1418 $\pm$ 0.0010	0.7711 $\pm$ 0.0029
	CRU	0.1138 $\pm$ 0.0018	0.0866 $\pm$ 0.0016	0.7326 $\pm$ 0.0086	0.1913 $\pm$ 0.0040	0.1236 $\pm$ 0.0018	0.8226 $\pm$ 0.0075	0.2021 $\pm$ 0.0013	0.1551 $\pm$ 0.0011	0.7375 $\pm$ 0.0034
	Mamba	0.0964 $\pm$ 0.0017	0.0706 $\pm$ 0.0015	0.8082 $\pm$ 0.0067	0.1715 $\pm$ 0.0044	0.1096 $\pm$ 0.0029	0.8573 $\pm$ 0.0074	0.2006 $\pm$ 0.0056	0.1502 $\pm$ 0.0053	0.7412 $\pm$ 0.0142
	ENC-ODE (ours)	0.0882$\pm$0.0023	0.0621$\pm$0.0006	0.8435$\pm$0.0044	0.1443$\pm$0.0013	0.0956$\pm$0.0004	0.8991$\pm$0.0017	0.1855$\pm$0.0028	0.1394$\pm$0.0021	0.7788$\pm$0.0067
Multimodal	RNN	0.1006 $\pm$ 0.0001	0.0768 $\pm$ 0.0001	0.7900 $\pm$ 0.0004	0.2401 $\pm$ 0.0030	0.1448 $\pm$ 0.0022	0.5367 $\pm$ 0.0122	0.2171 $\pm$ 0.0016	0.1597 $\pm$ 0.0009	0.6974 $\pm$ 0.0045
	Transformer	0.0984 $\pm$ 0.0010	0.0751 $\pm$ 0.0007	0.7991 $\pm$ 0.0037	0.2122 $\pm$ 0.0054	0.1271 $\pm$ 0.0027	0.6401 $\pm$ 0.0181	0.1995 $\pm$ 0.0014	0.1503 $\pm$ 0.0007	0.7444 $\pm$ 0.0037
	Neural ODE	0.0914 $\pm$ 0.0004	0.0694 $\pm$ 0.0003	0.8265 $\pm$ 0.0014	0.2671 $\pm$ 0.0045	0.1668 $\pm$ 0.0029	0.4302 $\pm$ 0.0192	0.2806 $\pm$ 0.0024	0.2108 $\pm$ 0.0019	0.4944 $\pm$ 0.0087
	ODE-RNN	0.0936 $\pm$ 0.0010	0.0715 $\pm$ 0.0008	0.8179 $\pm$ 0.0040	0.2276 $\pm$ 0.0023	0.1346 $\pm$ 0.0020	0.5862 $\pm$ 0.0085	0.2391 $\pm$ 0.0037	0.1784 $\pm$ 0.0031	0.6330 $\pm$ 0.0115
	Neural Flows	0.1155 $\pm$ 0.0018	0.0890 $\pm$ 0.0014	0.7230 $\pm$ 0.0085	0.2742 $\pm$ 0.0052	0.1842 $\pm$ 0.0045	0.3995 $\pm$ 0.0227	0.2234 $\pm$ 0.0024	0.1708 $\pm$ 0.0019	0.6795 $\pm$ 0.0069
	CRU	0.1157 $\pm$ 0.0024	0.0895 $\pm$ 0.0020	0.7217 $\pm$ 0.0117	0.2711 $\pm$ 0.0045	0.1836 $\pm$ 0.0033	0.4127 $\pm$ 0.0196	0.2305 $\pm$ 0.0026	0.1742 $\pm$ 0.0027	0.6590 $\pm$ 0.0076
	Mamba	0.0959 $\pm$ 0.0021	0.0729 $\pm$ 0.0015	0.8091 $\pm$ 0.0082	0.2016 $\pm$ 0.0046	0.1275 $\pm$ 0.0035	0.6752 $\pm$ 0.0147	0.2059 $\pm$ 0.0038	0.1536 $\pm$ 0.0021	0.7277 $\pm$ 0.0100
	ENC-ODE (ours)	0.0760$\pm$0.0006	0.0579$\pm$0.0005	0.8802$\pm$0.0019	0.1397$\pm$0.0002	0.0921$\pm$0.0005	0.8442$\pm$0.0005	0.1781$\pm$0.0022	0.1351$\pm$0.0019	0.7964$\pm$0.0050

epochs using Adam [16] optimizer with a learning rate of 1e-4, a batch size of 1024, and a weight decay of 0.05.

Event Prediction Settings. Given a sequence of observations made on continuous trajectory $f(t)$ for $t \in [t_1, t_p)$, our goal is to predict $f(t_p | \mathcal{H}_{t_{p-1}})$. Since observations occur at discrete time points, we optimize the model based on discrete data. During the training, the objective simplifies to predicting the final event $f(t_p)$ from the preceding $p - 1$ events. We conduct **1)** unimodal and **2)** multimodal sequence predictions to evaluate the contribution of individual biomarkers and the benefits of multimodal integration for contextual disease progression modeling. In the unimodal setting, sequences include only one target modality (e.g., AMY) with a minimum length of 2. In the multimodal setting, sequences contain at least one instance of the target modality before the final observation, which is also the target modality, while non-target modalities can appear at any position elongating the sequence as summarized in Tab. 2.

To evaluate our method in capturing the temporal dynamics of Alzheimer’s biomarkers, we compare its performance against representative sequence prediction models such as RNN [24], Transformer [27], Neural ODE [8], ODE-RNN [23], Neural Flow [5], CRU [25] and Mamba [13]. For fair comparisons, we designed baselines to have similar model capacity (i.e., the number of parameters). Also, we used the same set of training hyperparameters described in Sec. 3.1 across all baselines, and applied them consistently across different modalities. Root mean squared error (RMSE), mean absolute error (MAE) and coefficient of determination (R^2) were employed as evaluation metrics to quantify overall predictive performance. To ensure the robustness and reliability of our results, we replicated the same experiment five times and report the mean and standard deviation.




Fig. 2: Visualization of predicted FDG SUVR sequences (subject ID: 129_S_4422). The model captures the continuous decline in glucose metabolism as the disease progresses from CN to AD. Green boxes denote actual FDG observations at t_3 and t_8 .

Table 4: Ablation on diagnosis driven ODEs and attention aggregation using RMSE.

	FDG	Tau	AMY
w/o Diagnosis ODEs	0.0816 $\pm$ 0.006	0.1511 $\pm$ 0.041	0.1818 $\pm$ 0.012
Uniform Aggregation	0.0918 $\pm$ 0.003	0.2315 $\pm$ 0.007	0.2114 $\pm$ 0.004
Latest Observation Only	0.0988 $\pm$ 0.002	0.2377 $\pm$ 0.023	0.2800 $\pm$ 0.005
ENC-ODE (Ours)	0.0760$\pm$0.006	0.1397$\pm$0.001	0.1781$\pm$0.002

Table 5: Average training time per epoch with and without Eq.(4).

Sequential (w/o Eq.(4))	25.66s
Parallel (w/ Eq.(4))	10.47s
Time Ratio	40.8%

3.2 Experimental Results and Analysis

Comparison on Prediction Score. Tab. 3 demonstrates that ENC-ODE outperforms baselines in both unimodal and multimodal settings, achieving the highest performance across all metrics. In unimodal scenarios, its superior performance stems from the diagnosis-conditioned ODEs, which precisely characterize continuous biomarker evolution by accounting for stage-specific progression rates over irregular time intervals. Notably, ENC-ODE gains performance from incorporating different modality types in prediction, whereas others rather deteriorated in many cases. This highlights that target-aware attention effectively mitigates information interference by bypassing restrictive compression and selectively aggregating relevant features across disparate modalities.

Visualization of Actual Brain Trajectories. For a multimodal sequence of length 11 from a subject (ID: 129_S_4422), Fig. 2 illustrates the predicted evolution of FDG SUVR over time. We include ground truth values at t_3 and t_8 to provide a direct comparison at the specific time points where FDG were measured. As discussed in literature [19,22], FDG generally decreases over time and shows a noticeable decline when there is a significant change in the diagnostic label (e.g., EMCI to AD at t_9). Also, abrupt changes in parahippocampal gyrus and posterior cingulate show the clinical validity based on articles [18,1,15]. The close alignment between our predicted trajectories $\hat{f}_t$ and the ground truth measurements $f(t)$ at t_3 and t_8 further reinforces the reliability of our framework.

Ablation on accuracy. Tab. 4 summarizes the impact of our proposed modules, averaged over five different random seeds. To assess the influence of the diagnosis information in ODE modeling, we removed the diagnosis term $\xi_\delta(\delta_{t_i})$




Fig. 3: Left: Attention Score from different modalities, Right: Averaged rate of change during $t \in [15, 30]$ in FDG prediction from all modalities $m_t \in \{AMY, FDG, TAU\}$.

from Eq. (2) as follows: $dh(t|e_{t_i}) = \gamma(h(t|e_{t_i}); \theta)dt$. Furthermore, we evaluated the necessity of our attention mechanism by comparing it against uniform aggregation and a baseline using only the latest observation. The removal or simplification of each component led to a noticeable drop in performance across all settings, resulting in significantly higher RMSE values. In particular, simplifying attention module showed significant degradation in Tau prediction, where Tau measurements account for only 8.9% of the total observations (Tab. 1). These results demonstrate that relying on a single recent event or simple averaging is insufficient to capture the complex dependencies within multimodal histories.

Ablation on efficiency. Tab. 5 highlights the efficiency of our parallel training strategy (Eq. (4)). Unlike a sequential approach regarding each hidden state trajectory as an independent ODE, our method reduces the computational time, including backpropagation, by 59.2%, demonstrating its scalability and efficiency.

Attention Weight Analysis. The performance degradation under uniform aggregation or exclusive reliance on the latest observation (Tab. 4) underscores the necessity of our selective weighting strategy. While it is straightforward that the attention mostly focuses on the same modality type, to examine which modality provides complementary information in prediction, we visualize the relative attention scores of two other modality types in Fig. 3 (left). Given that FDG contributes significantly to the predictions, Fig. 3 (right) illustrates the temporal evolution of FDG-driven influences on each biomarker. Notable changes in Amyloid and Tau stand out in regions including inferior temporal, cuneus, and posterior cingulate cortex, which are established AD-specific regions [14, 9, 1].

4 Conclusion

In this work, we presented ENC-ODE, a continuous-time framework designed to navigate the temporal irregularity of multimodal neurodegenerative disease data. Our approach leverages diagnosis-conditioned dynamics to capture the nuanced, stage-specific tempo of brain biomarker evolution. In addition, the integration of target-aware attention enables a selective deconstruction of multimodal histories, effectively bypassing the information bottlenecks of traditional sequence models while providing a clear window into event-level influences. Validated on the ADNI dataset, ENC-ODE not only achieves superior predictive performance

but also offers a clinically grounded and computationally efficient paradigm for longitudinal analysis. This work provides a high-resolution predictive paradigm that can support personalized treatment planning through event-level analysis.

Acknowledgments. This research was supported by RS-2026-25494850 (60%), RS-2025-02216257 (35%), and RS-2019-II1091906 (AI Graduate Program at POSTECH, 5%).

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

References

1. Ali, D.G., Bahrani, A.A., et al.: Amyloid-pet levels in the precuneus and posterior cingulate cortices are associated with executive function scores in preclinical alzheimer’s disease prior to overt global amyloid positivity. *Journal of Alzheimer’s Disease* **88**(3), 1127–1135 (2022)
2. Baek, S., Choi, I., et al.: Learning covariance-based multi-scale representation of neuroimaging measures for alzheimer classification. In: ISBI. pp. 1–5. IEEE (2023)
3. Baek, S., Sim, J., et al.: Modality-agnostic style transfer for holistic feature imputation. In: ISBI. pp. 1–5. IEEE (2024)
4. Baek, S., Sim, J., et al.: Ocl: Ordinal contrastive learning for imputating features with progressive labels. In: MICCAI. pp. 334–344. Springer (2024)
5. Biloš, M., Sommer, J., Rangapuram, S.S., Januschowski, T., Günnemann, S.: Neural flows: Efficient alternative to neural odes. *NeurIPS* **34**, 21325–21337 (2021)
6. Biswas, B.N., Chatterjee, S., Mukherjee, S., Pal, S.: A discussion on euler method: A review. *Electronic Journal of Mathematical Analysis and Applications* **1**(2), 2090–2792 (2013)
7. Blennow, K., Zetterberg, H.: Biomarkers for alzheimer’s disease: current status and prospects for the future. *Journal of internal medicine* **284**(6), 643–663 (2018)
8. Chen, R.T., Rubanova, Y., Bettencourt, J., et al.: Neural ordinary differential equations. *NeurIPS* **31** (2018)
9. Cho, H., Seo, S.W., Kim, J.H., et al.: Amyloid deposition in early onset versus late onset alzheimer’s disease. *Journal of Alzheimer’s Disease* **35**(4), 813–821 (2013)
10. Cohen, A.D., Klunk, W.E.: Early detection of alzheimer’s disease using pib and fdg pet. *Neurobiology of disease* **72**, 117–122 (2014)
11. 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)
12. Goenka, N., Tiwari, S.: Deep learning for alzheimer prediction using brain biomarkers. *Artificial Intelligence Review* **54**(7), 4827–4871 (2021)
13. Gu, A., Dao, T.: Mamba: Linear-time sequence modeling with selective state spaces. In: COLM (2024)
14. Insel, P.S., Mormino, E.C., et al.: Neuroanatomical spread of amyloid β and tau in alzheimer’s disease: implications for primary prevention. *Brain Communications* **2**(1), fcaa007 (2020)
15. Kato, T., Inui, Y., Nakamura, A., et al.: Brain fluorodeoxyglucose (fdg) pet in dementia. *Ageing research reviews* **30**, 73–84 (2016)

16. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
17. Liu, H., Brock, A., Simonyan, K., Le, Q.: Evolving normalization-activation layers. *NeurIPS* **33**, 13539–13550 (2020)
18. Mosconi, L., Tsui, W.H., De Santi, S., et al.: Reduced hippocampal metabolism in mci and ad: automated fdg-pet image analysis. *Neurology* **64**(11), 1860–1867 (2005)
19. Mosconi, L., Mistur, R., Switalski, R., et al.: Fdg-pet changes in brain glucose metabolism from normal cognition to pathologically verified alzheimer’s disease. *European journal of nuclear medicine and molecular imaging* **36**, 811–822 (2009)
20. Mueller, S.G., Weiner, M.W., Thal, L.J., et al.: The alzheimer’s disease neuroimaging initiative. *Neuroimaging Clinics* **15**(4), 869–877 (2005)
21. Nedialkov, N.S., Jackson, K.R., Corliss, G.F.: Validated solutions of initial value problems for ordinary differential equations. *Applied Mathematics and Computation* **105**(1), 21–68 (1999)
22. Rice, L., Bisdas, S.: The diagnostic value of fdg and amyloid pet in alzheimer’s disease—a systematic review. *European journal of radiology* **94**, 16–24 (2017)
23. Rubanova, Y., Chen, R.T., Duvenaud, D.K.: Latent ordinary differential equations for irregularly-sampled time series. *NeurIPS* **32** (2019)
24. Rumelhart, D.E., et al.: Learning representations by back-propagating errors. *nature* **323**(6088), 533–536 (1986)
25. Schirmer, M., Eltayeb, M., Lessmann, S., Rudolph, M.: Modeling irregular time series with continuous recurrent units. In: *ICML*. pp. 19388–19405. PMLR (2022)
26. Thie, J.A.: Understanding the standardized uptake value, its methods, and implications for usage. *Journal of Nuclear Medicine* **45**(9), 1431–1434 (2004)
27. Vaswani, A., Shazeer, N., Parmar, N., et al.: Attention is all you need. *NeurIPS* **30** (2017)
28. Yang, C., Mei, H., Eisner, J.: Transformer embeddings of irregularly spaced events and their participants. In: *ICLR* (2022)