

DISENTANGLED REPRESENTATION OF LONGITUDINAL β -AMYLOID FOR AD VIA SEQUENTIAL GRAPH VARIATIONAL AUTOENCODER WITH SUPERVISION

Fan Yang^{*} Guorong Wu[†] Won Hwa Kim^{*‡}

^{*} University of Texas at Arlington, Arlington, USA

[†] University of North Carolina, Chapel Hill, Chapel Hill, USA

[‡] Pohang University of Science and Technology, Pohang, South Korea

ABSTRACT

The emergence of Positron Emission Tomography (PET) imaging allows us to quantify the burden of amyloid plaques *in-vivo*, which is one of the hallmarks of Alzheimer’s disease (AD). However, the invasive exposure to radiation and high imaging cost significantly restrict the application of PET in characterizing the evolution of pathology burden which often requires longitudinal PET image sequences. In this regard, we propose a proof-of-concept solution to generate the complete trajectory of pathological events throughout the brain based on very limited number of PET scans. We present a novel variational autoencoder model to learn a latent population-level representation of neurodegeneration process based on the longitudinal β -amyloid measurements at each brain region and longitudinal diagnostic stages. As the propagation of pathological burdens follow the topology of brain connectome, we further cast our neural network into a supervised sequential graph VAE, where we use the brain network to guide the representation learning. Experiments show that the disentangled representation can capture disease-related dynamics of amyloid and forecast the level of amyloid depositions at future time points.

1. INTRODUCTION

Alzheimer’s disease (AD) is a progressive neurodegenerative disorder resulting in cognitive impairments that interfere with functions for daily life [1]. In current pathophysiologic understanding of AD, β -amyloid deposition occurs as the first pathological event, followed by tau fibrillary tangles as the downstream effect, and chronologically lead to neurodegeneration as the final signal of cognitive disorders [2]. Thus, accumulation of β -amyloid neuritic plaques is one of the earliest risk factor for potential development of AD [3] and one of the key proteins to assess in understanding AD [2, 4].

Non-invasive Positron Emission Tomography (PET) provides a direct measure of *in vivo* β -amyloid status to better characterize early AD. Despite the improved diagnostic ability from amyloid PET, the cost for PET scans are very expensive ($\sim$ \$5,000) which prevents its widespread clinical adop-

tion. Under limited amyloid PET data at present, to facilitate clinical research on β -amyloid, a framework that can learn dynamics of amyloid and forecast future measures from limited past timestamps may be of great interest, helping understand how amyloid functions and its critical roles in AD.

Many studies have shown that structural brain connectivity from Diffusion Tensor Imaging (DTI) is highly associated with AD progression [5, 6, 7], whose topology between regions of interest (ROIs) behaves as a path for the amyloid deposition [8]. However, it is very challenging to learn complex dynamics of β -amyloid over brain networks through noisy data from limited timestamps per subject (e.g., less than 3 timestamps on average). Given limited and noisy longitudinal β -amyloid PET measures over a structural brain network, our aim is to develop a framework to uncover the disease dynamics or progression pattern of β -amyloid and forecast amyloid depositions at future timestamps to better characterize the progression of AD. Unfortunately, this is not trivial as observations from PET scans are complexly affected by several variables such as age, gender, anatomy, disease effect, etc.

Notice that these variables can be separated as time-varying and time-invariant components, where the progression of AD is a major factor for the time-varying one. Therefore, separating them in a latent space is critical; disentangled representation learning would be able to disentangle time-invariant contents (e.g., anatomical information) from time-varying contents (e.g., morphological changes, dynamics of amyloid). Such disentangled representations not only help model becomes more explainable, but also can benefit conditional data generation for downstream tasks [9, 10], e.g., cross-modality registration and segmentation [11, 12].

To this end, motivated by the schematics of Sequential Graph VAE [13, 14, 15], we develop a framework that learns a latent disentangled representation composed of time-varying and time-invariant latent components to characterize longitudinal β -amyloid over the structural brain network. The core idea is to capture disease-related dynamics of β -amyloid as well as forecast future amyloid depositions using the disentangled representation. The major contributions of this work are: 1) We incorporated “time-dependent” label as a super-




Fig. 1: A graphical visualisation of the encoder and decoder. (a) Encoder: time-invariant random variable (r.v.) f are inferred by time-dependent labels y and data X , and time-varying r.v. z are sequentially inferred by labels y and data X ; (b) Decoder: data are sequentially generated from time-dependent labels y , time-invariant r.v. f and time-varying r.v. z via latent r.v. W .

vision into the model to characterize longitudinal effect; 2) We integrated a brain network to make the framework more robust to the subject-wise heterogeneous dynamics; 3) We validated this framework on longitudinal β -amyloid over the brain network with diagnostic labels of AD from Alzheimer’s Disease Neuroimaging Initiative (ADNI). Experimental results suggest a significant potential that this framework will facilitate clinical research by enriching amyloid data collection and help better understand the role of amyloid in the progression of AD before the disease symptoms manifest.

Related Work. Disentangled representation learning aims to map original data onto a low-dimensional space to help extract semantically meaningful components. β -VAE and follow-up works [16, 17, 18] proposed a heavier penalty on KL divergence term in objective function to learn a better disentanglement representation by encouraging learning of independent factors. FHVAE [14] and DSVAE [15] were developed for sequential data generation by elaborately designing model architectures and disentangling latent factors into time-invariant and time-varying parts. However, less research has been conducted for disentangled representations of *sequences* with the capacity which incorporates time-dependent labels and graph structure and performs forecasting as well. Furthermore, despite several works on disease progressing exist [19, 20, 21, 22], representation disentanglement of medical sequential data has been under-explored.

2. METHODS

2.1. Supervised Sequential Graph VAE Model

Supervised Sequential Graph Variational Autoencoder (SSG-VAE) is designed for learning a disentangled representation composed of time-variant and time-invariant latent variables to capture the dynamics of longitudinal measures and forecast the measure at the future timestamp. We observed sequential data that appear as M^{sup} pairs of supervised data points $\mathcal{D}^{sup} = \{X_i, y_i\}_{i=1}^{M^{sup}}$ over a shared graph G with N nodes, where $X_i = (X_{i,1}, \dots, X_{i,T_i})$ refer to the i -th sequential observations, i.e., $X_{i,t} \in \mathbb{R}^N$, and $y_i = (y_{i,1}, \dots, y_{i,T_i})$ denote the corresponding time-dependent diagnostic labels. We

will leave out the index i wherever it is clear that the terms we are referring to are associated with a single data point. SSG-VAE simultaneously trains a probabilistic encoder and decoder, and factorizes latent variables into two disentangled variables: time-invariant variable f and time-varying variable $z_{1:T} = (z_1, \dots, z_T)$. We expect that latent variable f can encode the time-invariant global aspects of the data, while latent variable z will encode how the time-varying information at timestamp t is morphed into that of timestamp $t+1$. The architectures of decoder and encoder are visualised in Figure 1.

2.1.1. Objective Function

We design a supervised variational autoencoder framework with an objective function [23, 24] defined over $\mathcal{D}^{sup}$ as

$$\begin{aligned} \mathcal{L}^{sup}(\theta, \phi; \mathcal{D}^{sup}) \\ = \mathbb{E}_{\hat{p}(\mathbf{X}, \mathbf{y})} \left[\mathbb{E}_{q_{\phi}(f, \mathbf{z} | \mathbf{X}, \mathbf{y})} \left[\ln \frac{p_{\theta}(\mathbf{X}, \mathbf{y}, f, \mathbf{z})}{q_{\phi}(f, \mathbf{z} | \mathbf{X}, \mathbf{y})} \right] \right], \end{aligned} \quad (1)$$

where the model parameters of the decoder is denoted as θ and the model parameters of the encoder is denoted as ϕ . Here, $\hat{p}(\cdot)$ denotes the empirical distribution, $p_{\theta}(\cdot)$ refers to the decoder distribution and $q_{\phi}(\cdot)$ is the variational posterior.

2.1.2. Prior

The prior of time-invariant latent variable f is defined as $f \sim \mathcal{N}(\mathbf{0}, I)$. We assume that time-varying latent variables $z_{1:T}$ follow a sequential prior $z_t | z_{t-1} \sim \mathcal{N}(\mu_t, \text{diag}(\sigma_t^2))$, where $[\mu_t, \sigma_t]$ are the parameters of the prior distribution and is parameterized as a recurrent network LSTM [25], in which the hidden states are updated temporally. Moreover, prior distributions of time-dependent labels y follow the multinomial distribution, i.e., $p_{\theta}(y_t) = \text{Cat}(y_t | \pi)$. Assuming labels y , time-invariant f and time-varying $z_{1:T}$ are mutually independent, the joint prior $p_{\theta}(y, f, z_{1:T})$ can be factorized as

$$p_{\theta}(y, f, z_{1:T}) = \prod_{t=1}^T p_{\theta}(y_t) p_{\theta}(f) \prod_{t=1}^T p_{\theta}(z_t | z_{<t}). \quad (2)$$

We use independent priors to regularize latent variables to be as independent as possible.

2.1.3. Generative Model: Decoder

The generative model is formalized by the factorization as

$$p_{\theta}(\mathbf{X}, \mathbf{y}, \mathbf{f}, \mathbf{z}_{1:T}) = \prod_{t=1}^T p_{\theta}(y_t) p_{\theta}(\mathbf{f}) \prod_{t=1}^T p_{\theta}(\mathbf{z}_t | \mathbf{z}_{<t}) p_{\theta}(W_t | y_t, \mathbf{f}, \mathbf{z}_t) p_{\theta}(X_t | W_t), \quad (3)$$

where $\mathbf{W} = (W_1, \dots, W_T)$ denotes the P -dimensional latent vectors and $W_t \in \mathbb{R}^{N \times P}$. Latent vectors $\mathbf{W}$ are generated from the time-dependent $\mathbf{y}$ and the two disentangled variables, i.e., the time-invariant $\mathbf{f}$ and the time-varying $\mathbf{z}$. We assume that data sequences $\mathbf{X}$ are generated from latent vectors $\mathbf{W}$ via the graph convolution as

$$X_t = \hat{A} W_t \Theta, \quad (4)$$

where Θ is the trainable weight matrix, A is the adjacent matrix of the graph G , $\hat{A} = \tilde{D}^{-1/2} \tilde{A} \tilde{D}^{-1/2}$, $\tilde{A} = A + I$, and $\tilde{D}_{ii} = \sum_j \tilde{A}_{ij}$. That is, we incorporate the topology of the graph G into the generative process using graph convolution.

2.1.4. Inference Model: Encoder

We use variational model $q_{\phi}(\mathbf{f}, \mathbf{z}_{1:T} | \mathbf{X}, \mathbf{y})$ to approximate the true posterior distribution $p_{\theta}(\mathbf{f}, \mathbf{z}_{1:T} | \mathbf{X}, \mathbf{y})$ over latent variables given data [26, 27]. The inference is factorized as

$$q_{\phi}(\mathbf{f}, \mathbf{z}_{1:T} | \mathbf{X}, \mathbf{y}) = q_{\phi}(\mathbf{f} | X_{1:T}, y_{1:T}) \prod_{t=1}^T q_{\phi}(\mathbf{z}_t | X_{\leq t}, y_{\leq t}). \quad (5)$$

The time-invariant $\mathbf{f}$ are conditional on the entire time sequences of data and time-dependent $\mathbf{y}$, while the time-dependent $\mathbf{z}_t$ is inferred from the sequences before timestamp t , i.e., $X_{\leq t}$ and $y_{\leq t}$. We model both $\mathbf{f}$ and $\mathbf{z}$ via LSTM.

2.2. Predictive Supervised Sequential Graph VAE Model

Although SSG-VAE model in Section 2.1 can successfully learn a disentangled representation which decomposes the static, time-varying and label information, it cannot be utilized for data forecasting at future timestamps, due to lack of the forecasting layer in the model. To conquer this challenge, we extend our SSG-VAE to Predictive Supervised Sequential Graph VAE (PSSG-VAE), which allows us to forecast latent variables and output outcomes at future time stamps.

Assuming the complete data pairs $\mathbf{X} = (X_1, \dots, X_T)$ and $\mathbf{y} = (y_1, \dots, y_T)$, we denote the observed data pairs as $\tilde{\mathbf{X}} = (X_1, \dots, X_{T-1})$ and $\tilde{\mathbf{y}} = (y_1, \dots, y_{T-1})$, with X_T, y_T denotes the forecast data pair. Then we can reformulate the inference model in Section 2.1.4 as

$$q_{\phi}(\mathbf{f}, \mathbf{z}_{1:T} | \mathbf{X}, \mathbf{y}) = q_{\phi}(\mathbf{f}, \mathbf{z}_{1:T-1} | \tilde{\mathbf{X}}, \tilde{\mathbf{y}}) q_{\phi}(\mathbf{z}_T | \mathbf{z}_{T-1}), \quad (6)$$

where $q_{\phi}(\mathbf{z}_T | \mathbf{z}_{T-1})$ can be any parametric function and we use a naive linear transition model for the rest of this work, and $q_{\phi}(\mathbf{f}, \mathbf{z}_{1:T-1} | \tilde{\mathbf{X}}, \tilde{\mathbf{y}})$ can be factorized similarly as Eq. 5. With aforementioned reformulation, our extended PSSG-VAE will be capable of data forecasting at future timestamps.

3. EXPERIMENTAL RESULTS

3.1. ADNI Dataset and Experimental Setup

Total of $N=720$ subjects were taken from the ADNI study that contained both amyloid PET and DTI images. Longitudinal β -amyloid data were processed from amyloid PET scans, and structural connectivity matrices (i.e., number of fiber tracts connecting different ROIs) were derived from DTI registered at Destrieux atlas in FreeSurfer [28] using a in-house tractography pipeline. Specifically, standardized uptake value ratio (SUVR) was computed for β -amyloid at each brain region, and an overall graph was obtained by taking the average of connectivity matrices from healthy subjects. Diagnostic labels of each scan categorize subjects' dementia stage as one of Cognitive Normal (CN), Significant Memory Concern (SMC), Early Mild Cognitive Impairment (EMCI), Late Mild Cognitive Impairment (LMCI) and Alzheimer's Disease (AD). Demographics of the subjects are presented in Table 1.

Table 1: Demographics of the ADNI dataset.

Demographics	CN	SMC	EMCI	LMCI	AD
# of Subjects	204	80	240	107	89
Gender (M/F)	106:98	27:53	138:102	61:46	44:45
Age (mean,std)	73.33(6.06)	71.06(5.01)	70.91(7.16)	72.04(7.92)	73.22(7.48)

CN: Cognitive Normal; SMC: Significant Memory Concern; EMCI: Early Mild Cognitive Impairment; LMCI: Late Mild Cognitive Impairment; AD: Alzheimer's Disease.

We validated our framework on the dataset from ADNI study including longitudinal β -amyloid data on brain networks with diagnostic stage labels of AD. We conducted three experiments as described below. Section 3.2 displays latent traversals over labels, where we explored how patterns of generated amyloid will change corresponding to variations of diagnostic labels. Section 3.3 shows reconstruction performance on the dynamics of β -amyloid, compared with the ground truth and visualized on brain surfaces. Section 3.4 demonstrates forecasting performances at the future timestamp with 3-fold cross validation. Here, since each subject has different number of visits, we propose two baseline approaches for comparison with our approach. One is the averaging where the amyloid measure at the last timestamp is estimated as the average of all historical measures at previous timestamps, the other is the linear regression by leveraging the average from all past timestamps as predictors. Root mean square error (RMSE) and Mean Absolute Error (MAE) between the ground truth and the predicted are used as the metrics for evaluation of forecasting performance.

3.2. Latent Traversals over Labels

To explore the relation between labels and generated amyloid, we conducted latent traversals task over labels of diagnostic stages using proposed SSG-VAE in Section 2.1. Specifically, we fixed a time-invariant variable $\mathbf{f}$ and a time-varying vari-




Fig. 2: Latent Traversals over Labels. From left to right: generated β -amyloid on brain surfaces with latent variables fixed and diagnostic labels varying from CN to AD, respectively. Label-related patterns match with existing knowledge from the neuroscience domain.

able z , and we varied the corresponding label from CN to AD respectively. We reconstructed the β -amyloid with those different labels but the same other latent variables shown in Figure 2. It is clear that as the status of disease stage becomes worse, the values of amyloid measure become larger, matching the existing knowledge from the neuroscience domain.

3.3. Reconstruction on the Dynamics of β -Amyloid

Here we illustrate that our SSG-VAE model can learn the complex dynamics of β -amyloid by showing the reconstruction results on the testing data. We show the true amyloid and reconstructed amyloid on brain surfaces in Figure 3. It visually demonstrates that the reconstruction not only captures the anatomical information but also successfully learns the true dynamics from the limited and noisy longitudinal data, as it resembles the patterns of amyloid on brain surfaces across timestamps (see the color changes in Figure 3).




Fig. 3: Top: true brain surfaces for a randomly selected subject at timestamp t_0 , t_1 and t_2 (True). Bottom: reconstructed brain surfaces for the same subject at timestamp t_0 , t_1 and t_2 (Recon). Generated via BrainPainter [29].

3.4. Forecasting β -Amyloid at the Future Timestamp

We validated proposed PSSG-VAE in Section 2.2 on the forecasting task, and showed that the overall RMSEs on testing for average approach, regression approach and ours are 0.38, 0.22 and 0.19, respectively. Our approach attains the lowest overall RMSE. We also summarised RMSEs and MAEs across the diagnostic labels in Table 2. It indicates that our approach performs significantly better than regression approach at the earliest CN stage, which shows the advantage of our




Fig. 4: Boxplot of forecasting performance at future timestamp visualizing RMSEs for average (AVE), regression (LR) and our approach. Ours yields lowest overall RMSE and smaller variation compared to other approaches.

approach by capturing the earliest sign of cognitive decline at the preclinical stage. Moreover, we visualized RMSEs for average approach, regression approach and ours in Figure 4. It shows that the forecasts from our approach is more robust and has smaller variation, as demonstrated by the smaller interquartile range and fewer extreme outliers in the boxplot.

Table 2: Forecasting Performance across diagnostic stages.

		CN	SMC	EMCI	LMCI	AD
Average	RMSE	0.42	0.39	0.35	0.37	0.36
	MAE	0.32	0.31	0.28	0.30	0.28
Regression	RMSE	0.30	0.20	0.17	0.17	0.20
	MAE	0.21	0.16	0.13	0.13	0.15
Ours	RMSE	0.20	0.20	0.17	0.15	0.19
	MAE	0.15	0.15	0.13	0.11	0.14

CN: Cognitive Normal; SMC: Significant Memory Concern; EMCI: Early Mild Cognitive Impairment; LMCI: Late Mild Cognitive Impairment; AD: Alzheimer’s Disease.

4. CONCLUSION

Understanding dynamics of β -amyloid and forecasting future amyloid depositions will facilitate clinical research in the pre-clinical stages of AD. Here we developed a novel Supervised Sequential Graph VAE to learn a latent representation comprising time-varying and time-invariant information to characterize longitudinal β -amyloid over structural brain network. With the learned disentangled representation, our framework can capture the robust dynamics of amyloid and forecast future amyloid depositions from a few past time points.

Acknowledgement. Supported by GAANN Doctoral Fellowships in CSE at UTA, NSF IIS CRII 1948510, NIH RF1 AG059312, NIH R03 AG070701, IITP-2019-0-01906 funded by MSIT (AI Graduate Program at POSTECH) and Korea R&D Special Economic Zone Foundation grant funded by the Ministry of Science and ICT (Project ID : 1711150077).

5. REFERENCES

- [1] Mark S Henry et al., “The development of effective biomarkers for alzheimer’s disease: a review,” *IJGP*, vol. 28, no. 4, pp. 331–340, 2013.
- [2] Randall J Bateman et al., “Clinical and biomarker changes in dominantly inherited alzheimer’s disease,” *N Engl J Med*, vol. 367, pp. 795–804, 2012.
- [3] Clifford R Jack Jr et al., “Nia-aa research framework: toward a biological definition of alzheimer’s disease,” *A&D*, vol. 14, no. 4, pp. 535–562, 2018.
- [4] Joseph Theriault et al., “Determining amyloid- β positivity using 18f-azd4694 pet imaging,” *JNM*, vol. 62, no. 2, pp. 247–252, 2021.
- [5] Jacob W Vogel et al., “Spread of pathological tau proteins through communicating neurons in human alzheimer’s disease,” *Nature communications*, vol. 11, no. 1, pp. 1–15, 2020.
- [6] Ashish Raj et al., “A network diffusion model of disease progression in dementia,” *Neuron*, vol. 73, no. 6, pp. 1204–1215, 2012.
- [7] Xin Ma et al., “Learning multi-resolution graph edge embedding for discovering brain network dysfunction in neurological disorders,” in *IPMI*. Springer, 2021, pp. 253–266.
- [8] Seong Jae Hwang et al., “Associations between positron emission tomography amyloid pathology and diffusion tensor imaging brain connectivity in pre-clinical alzheimer’s disease,” *Brain connectivity*, vol. 9, no. 2, pp. 162–173, 2019.
- [9] Yizhe Zhu et al., “S3vae: Self-supervised sequential vae for representation disentanglement and data generation,” in *CVPR*, 2020, pp. 6538–6547.
- [10] Jiahong Ouyang, Ehsan Adeli, et al., “Representation disentanglement for multi-modal mr analysis,” *arXiv preprint arXiv:2102.11456*, 2021.
- [11] Chen Qin, Bibo Shi, et al., “Unsupervised deformable registration for multi-modal images via disentangled representations,” in *IPMI*. Springer, 2019, pp. 249–261.
- [12] Junlin Yang et al., “Cross-modality segmentation by self-supervised semantic alignment in disentangled content space,” in *DART and DCL*. Springer, 2020.
- [13] Fan Yang et al., “Disentangled sequential graph autoencoder for preclinical Alzheimer’s disease characterizations from adni study,” in *MICCAI*, 2021, pp. 362–372.
- [14] Wei-Ning Hsu et al., “Unsupervised learning of disentangled and interpretable representations from sequential data,” in *NIPS*, 2017, pp. 1878–1889.
- [15] Li Yingzhen et al., “Disentangled sequential autoencoder,” in *ICML*. PMLR, 2018, pp. 5670–5679.
- [16] Irina Higgins, Loïc Matthey, et al., “beta-vae: Learning basic visual concepts with a constrained variational framework,” in *ICLR*, 2017.
- [17] Ricky T. Q. Chen, Xuechen Li, et al., “Isolating sources of disentanglement in variational autoencoders,” in *NIPS*. 2018, vol. 31, Curran Associates, Inc.
- [18] Hyunjik Kim and Andriy Mnih, “Disentangling by factorising,” in *ICML*. PMLR, 2018, pp. 2649–2658.
- [19] Hubert M Fonteijn et al., “An event-based model for disease progression and its application in familial alzheimer’s disease and huntington’s disease,” *NeuroImage*, vol. 60, no. 3, pp. 1880–1889, 2012.
- [20] Bruno M Jedynak et al., “A computational neurodegenerative disease progression score: method and results with the alzheimer’s disease neuroimaging initiative cohort,” *Neuroimage*, vol. 63, no. 3, pp. 1478–1486, 2012.
- [21] Braden C Soper et al., “A hidden markov model for population-level cervical cancer screening data,” *Statistics in Medicine*, vol. 39, no. 25, pp. 3569–3590, 2020.
- [22] Rui Meng et al., “Hierarchical hidden markov jump processes for cancer screening modeling,” *arXiv preprint arXiv:1910.05847*, 2019.
- [23] DP. Kingma et al., “Semi-supervised learning with deep generative models,” in *NIPS*, 2014, pp. 3581–3589.
- [24] Siddharth N et al., “Learning disentangled representations with semi-supervised deep generative models,” in *NIPS*. 2017, vol. 30, Curran Associates, Inc.
- [25] Sepp Hochreiter et al., “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [26] Michael I Jordan et al., “An introduction to variational methods for graphical models,” *Machine learning*, vol. 37, no. 2, pp. 183–233, 1999.
- [27] DP. Kingma and M. Welling, “Auto-encoding variational bayes,” *CoRR*, vol. abs/1312.6114, 2014.
- [28] C. Destrieux et al., “Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature,” *Neuroimage*, vol. 53, no. 1, pp. 1–15, 2010.
- [29] Răzvan V Marinescu et al., “Brainpainter: A software for the visualisation of brain structures, biomarkers and associated pathological processes,” in *MBIA and MFCA*, pp. 112–120. Springer, 2019.