mmNormVAE: Normative Modeling on Multimodal Neuroimaging Data using Variational Autoencoders

Sayantan Kumar
Department of Computer Science
Washington University in St. Louis
St. Louis, MO 63130
sayantan.kumar@wustl.edu

Philip R. O. Payne
Department of Computer Science
Washington University in St. Louis
St. Louis, MO 63130
prpayne@wustl.edu

Aristeidis Sotiras
Department Radiology
Washington University in St. Louis
St. Louis, MO 63130
aristeidis.sotiras@wustl.edu

Abstract

Normative modelling is a popular method for studying brain disorders like Alzheimer's Disease (AD) where the normal brain patterns of cognitively normal subjects are modelled and can be used at subject-level to detect deviations relating to disease pathology. So far, deep learning-based normative frameworks have largely been applied on a single imaging modality. We aim to design a multi-modal normative modelling framework based on multimodal variational autoencoders (mmNormVAE) where disease abnormality is aggregated across multiple neuroimaging modalities (T1-weighted and T2-weighted MRI) and subsequently used to estimate subject-level neuroanatomical deviations due to AD.

1 Introduction

Normative modelling is a method to study heterogeneity in clinical cohorts, while allowing predictions at an individual subject. Assuming that neurodegenerative cohorts like Alzheimer's Disease (AD) manifest as outliers (anomalies) compared to a healthy distribution, normative models learn the characteristics of the healthy brain and quantify how each AD subject deviates from the normative (healthy controls) distribution level [4, 7]. Typically, a normative analysis learns separate models for each brain feature e.g using Gaussian Process Regression (GPR) [8], support vector regression [7]. Recently generative models like Variational Autoencoders (VAE) [6, 1], Adversarial autoencoders (AAE) [12] have been proposed to learn the complex non-linear interaction between features. These methods have a unimodal structure with a single encoder and decoder and can handle only one imaging modality. However, brain disorders like AD are multifactorial, showing deviations from the norm in features of multiple imaging modalities. Since each modality is sensitive to disease effects to a varying degree, it is important to develop normative models that can handle multiple modalities.

Contributions: In this work, we design a multi-modal variational autoencoder based normative modeling framework (mmNormVAE) that aggregates information from multiple modalities in the latent space, modeling a joint distribution between the different modalities. As a normative model, mmNormVAE learns the brain volume patterns of healthy subjects and is subsequently applied on AD subjects at test time. During inference, the joint latent space can be used to estimate modality-specific reconstructions and subsequently identify the brain regions with abnormal (statistically significant) deviations relating to AD pathology. We hypothesize that deviation maps generated by mmNormVAE

1st Workshop on Deep Generative Models for Health at NeurIPS 2023.

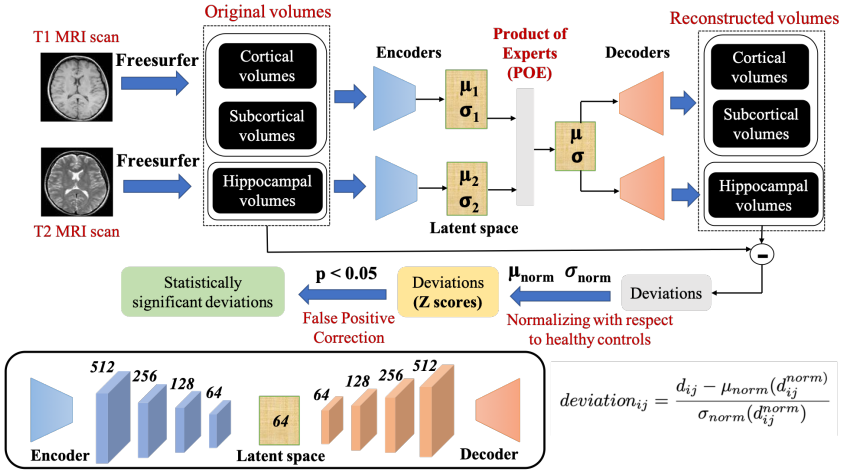

Figure 1: Our proposed multimodal normative modeling framework.The encoder and decoder networks have 4 fully-connected layers of sizes 512, 256, 128, 64 and 64, 128, 256, 512, respectively with a latent dimension of 64. The latent space parameters of the individual modalities were combined by the Product of Experts (POE) layer [3, 14] to form the shared latent space, which feeds the modality-specific decoders for reconstructions. The deviations in brain volumes of disease patients were calculated, normalized with respect to the healthy controls to form Z-scores. The regions with statistically significant deviations were identified ($p_{FDR} < 0.05$) after correcting for false positives (False Discovery Rate).

are more sensitive to disease staging, have a better correlation with patient cognition and estimate more regions with abnormal deviations compared to a unimodal baseline model.

2 Proposed Methodology

2.1 Multi-modal VAE

Let $X = [x_1, x_2, ...x_N]$ be a set of conditionally independent N modalities. The backbone of our model is a multimodal variational autoencoder (mVAE), a generative model of the form $p_\theta(x_1, x_2, ...x_N, z) = p(z)\prod_{i=1}^{N} p_\theta(x_i|z)$, where z is the latent variable and $p(z)$ is the prior. mVAE optimizes the ELBO (Evidence Lower Bound) which is a combination of modality-specific likelihood distributions $p_\theta\left(x_i|z\right)$ and the KL divergence between the approximate joint posterior $q(z|X)$ and prior $p(z)$ (Equation 1).

In mVAE, each modality can be treated as an "expert" and the approximate joint posterior can be estimated by taking the product of the unimodal posteriors $q_\phi(z|x_i)$ (Product-of-Experts (PoE)) (Equation 2) [14]. Assuming each encoder network (unimodal posterior) $q(z|x_i)$ follows a Gaussian distribution, the joint approimate posterior is also Gaussian with mean and variance represented by Equation 3 and 4 respectively [3].

$$\text{ELBO}(X) = \mathbb{E}_{q(z|X)}\left[\sum_{x_i \in X} \log p_\theta\left(x_i|z\right)\right] - \text{KL}\left[q_\phi(z|X)\|p(z)\right] \tag{1}$$

$$q_{PoE}(z \mid X) = p(z)\prod_{i=1}^{N} q_\phi(z \mid x_i) \tag{2}$$

$$\mu = (\sum_i \mu_i * T_i)(\sum_i T_i)^{-1} \tag{3}$$

$$\sigma = (\sum_i T_i)^{-1} \tag{4}$$

2.2 Multi-modal normative modelling

Cortical and subcortical brain volumes extracted from T1-weighted Magnetic Resonance Imaging (MRI) scans and hippocampal volumes extracted from T2-weighted MRI scans were used as input to two modality specific encoders. Information from multiple modalities is aggregated in the shared latent space to estimate the joint posterior which models the joint distribution between the 2 modalities. The shared latent parameters were fed through the modality-specific decoders to get the modality-specific reconstructions. Our multi-modal framework is first trained on cognitively unimpaired subjects to characterize the healthy cohort and subsequently applied on the disease cohort at test time. The main idea of the normative approach is that since model only learns how to reconstruct the brain region volumes of HC subjects, it will be less precise (more reconstruction error) in reconstructing the brain volumes of AD patients.

2.3 Calculating the deviations

For each disease patient, the deviations d_{ij} with respect to the healthy controls were calculated as the absolute signed difference between the original and reconstructed brain region volumes. Assuming that the proposed model is not able to perfectly reconstruct the brain region volumes of the healthy subjects, the deviations were normalized with respect to the mean μ_{norm} and variance σ_{norm} of deviations dij^{norm} of the healthy participants calculated from a separate held-out validation cohort. The final normalized deviation $deviation_{ij}$ (Z-scores) were calculated by normalizing the deviations of disease patients with respect to the mean μ_j and variance σ_j of the deviations of healthy participants for each brain region j. $deviation_{ij} = \frac{d_{ij} - \mu_{norm}(d_{ij}^{norm})}{\sigma_{norm}(d_{ij}^{norm})}$

We identified the brain regions of each patient whose deviations (Z-scores) were significantly different from those of the healthy control subjects ($p < 0.05$). Since the normalized deviations $deviation_{ij}$ were estimated independently for each brain region for every patient, FDR (False Discovery Rate) correction was applied to control the Type 1 error rate (false positive correction) [2]. Our proposed framework has been summarized in Figure 1.

3 Experimental Analysis

3.1 Data

For training the model framework, we selected 9633 healthy controls (HC) from the UKBiobank dataset [13] after excluding all subjects with recent history of anxiety, depression and nerve disorders. 862 disease patients (106 Significant Memory Concerns (SMC), 312 Early Mild Cognitive Impairment (EMCI), 263 Late Mild Cognitive Impairment (LMCI) and 181 AD) were selected from the Alzheimer's Disease Neuroimaging (ADNI) dataset [11]. For data harmonization between UKB and ADNI participants, the normative model estimated on UKB was re-trained on ADNI by a transfer learning approach. As patients progress from SMC to the AD stage, the severity of impairment increases. For both the datasets, we used the FreeSurfer software (version 5.1) [5] to estimate the volumes of 64 cortical and 35 subcortical brain regions from T1-weighted MRI images as well as the volumes of 16 hippocampal regions from T2-weighted MRI images, respectively.

3.2 Feature preprocessing

As part of the feature pre-processing step, the brain region volumes of each subject were normalized by their Intracranial Volume (ICV). The healthy controls from UK Biobank data was split into 80% for training and 20% for validation, which was used for early stopping to prevent overfitting. The 269 CN participants in ADNI were split into 70% for model training and validation, a 15% held-out validation set for estimating the parameters of the normative population and 15% in the test for estimating the deviations (see next subsection). The volumes of each region were scaled between 0 and 1 using MinMax scaling. The mean and variance calculated in the training set were also used to scale the data in validation and test sets.

3.3 Baselines and Training details

We compared our proposed framework with a baseline VAE having a unimodal VAE (single encoder and decoder) which takes the cortical, subcortical and hippocampal region volumes into a single concatenated input. We conditioned both the proposed and baseline VAE networks on the age and sex of patients, represented as one-hot encoding vectors, to ensure that the deviations in regional brain volumes reflect only the disease pathology and not deviations due to effects of covariates. Both the VAE models were trained using the Adam optimizer with model hyperparameters as follows: epochs = 500, learning rate = 10^-5, batch size = 256 and latent dimension = 64. The encoder and decoder networks have 4 dense layers of sizes $512, 256, 128, 64$ and $64, 128, 256, 512$ respectively.

4 Results

4.1 Sensitivity of deviation maps towards disease staging

The mean deviations of each patient across all 115 brain regions reflect the measure of abnormality or neuroanatomical alteration in the brain due to the AD progression (SMC $\rightarrow$ EMCI $\rightarrow$ LMCI $\rightarrow$ AD) and should ideally include more regions with increasing severity of the disease stages. For both our proposed multimodal and baseline unimodal framework, the patients exhibited more abnormality (mean deviations) with increasing severity of their condition from SMC to AD (**Figure 2 left**). Higher slope suggests that our proposed model is more sensitive to disease staging compared to the baseline. The pairwise differences in deviations between the disease categories were statistically significant ($p_{FDR} < 0.05$) except for the SMC and EMCI pair (**Figure 2 middle**). Thus, the deviations generated by our proposed model can better capture the neuroanatomical alterations in the brain due to the progressive stages of AD.

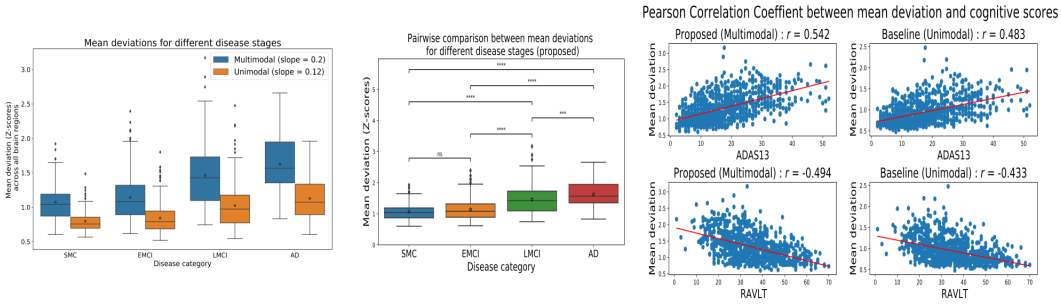

Figure 2: **Left**: Box plot showing the mean deviation of each patient (across all 115 brain regions) all four AD disease categories. The slope values shown in the figure legend are obtained by fitting a linear model across the mean value of each category (shown by * within each box plot). **middle**: Statistical annotations on the deviation maps generated by our proposed framework, showing the pairwise statistical comparison ($p_{FDR} < 0.05$) between the deviation maps in each disease category. ($ns : 0.05 < p <= 1; * : 0.01 < p <= 0.05; ** : 0.001 < p <= 0.01; * * * : 0.0001 < p <= 0.001; * * ** : p <= 0.0001$) **right**: Pearson Correlation between mean deviations and patient cognition represented by ADAS13 (top row) and RAVLT (bottom row). Each point in the plot represents a patient and the dark red line denotes the linear regression fit of the points.

4.2 Correlation of deviation maps with patient cognition

We analyzed the Pearson correlation coefficient between the patient-level deviations (mean deviatios across all 115 regions) and cognitive assessment scores, AD Assessment Scale (ADAS13) [10] and Rey Auditory Verbal Learning Test (RAVLT) [9]. ADAS13 test is a series of 13 cognitive tasks that can be used to assess the level of cognitive dysfunction in AD. RAVLT is a neuropsychological assessment to evaluate the nature and severity of memory impairment over time. High scores of ADAS13 and low scores of RAVLT indicate greater loss in memory and cognition. Our proposed framework exhibited higher correlation with both scores (high r value), compared to the baseline ($p < 0.05$) (**Figure 2 right**).

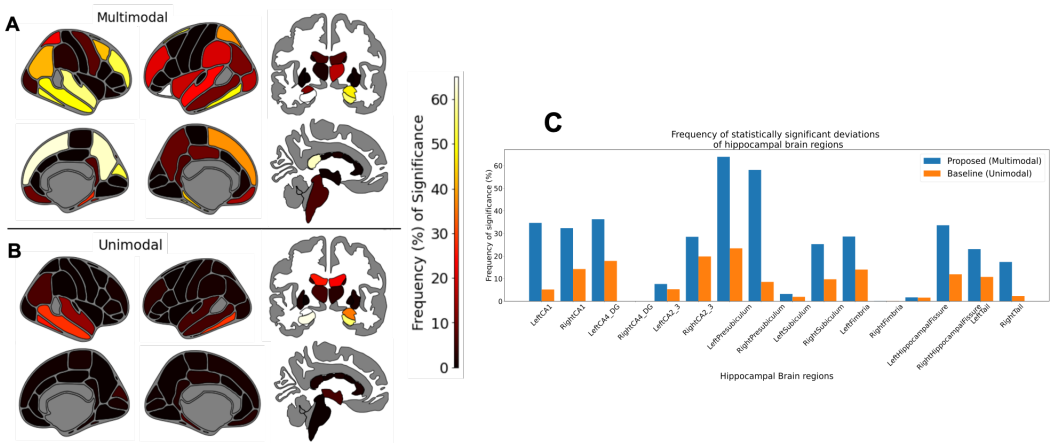

Figure 3: Frequency of significance: Number of times (%) each cortical and subcortical region (Figure 3A, 3B) and hippocampal region (Figure 3C) exhibited statistically significant deviations ($p_{FDR} < 0.05$) across all 862 patients in the test set. The 1st 2 columns in Figure 3A, 3B represent the 4 views of the Desikan-Killany cortical atlas and the last column represent the coronal and saggital views of the Aseg subcortical atlas. The colormap shows the frequency of significance in % with lighter colours representing higher frequency. The regions with higher frequency of significance across both the models (proposed multimodal and baseline unimodal) indicate the regions with abnormal brain deviation patterns due to AD.

4.3 Identifying abnormal brain deviations

We identified the brain regions whose deviations in brain volumes compared to the healthy controls were statistically significant after false positive correction ($p_{FDR} < 0.05$). For both the T1-weighted MRI and T2-weighted MRI modalities, we calculated how many times in % each cortical, subcortical region (Figure 3A, B) and hippocampal region (Figure 3C) had statistically significant deviations across all 862 patients in the ADNI test set. The deviations generated by our proposed model has higher fraction of significance for all the regions compared to the unimodal baseline. The regions with higher frequency of significance across both the models (proposed and baseline) indicate the regions with abnormal brain structural patterns. In other words, we can get an idea which of the brain regions have more deviation in volumes due to the progressive stages of AD compared to the other regions. Some of the regions with showing deviations include (i) **cortical** : *entorhinal-left, lateralorbitofrontal-left, fusiform-left, precuneus-right, superiorfrontal-right*, (ii) **subcortical** : *hippocampus-left*, *hippocampus-right*, *amygdala-right*, *Corpus-Callosum-Posterior* and (iii) **hippocampal** : *presubiculum-left*, *presubiculum-right*.

5 Conclusion and Future Work

We propose a multimodal VAE-based normative modelling framework (mmNormVAE), which models the joint distribution between brain region volumes derived from multiple MRI modalities (T1 and T2 MRI). Our framework quantifies at a subject-level how patients with different stages of AD deviate from the expected pattern learned from the healthy controls. The deviations generated by our proposed multimodal framework were more sensitive to disease staging within AD, had a better correlation with patient cognition and more regions with statistically significant deviations compared to a unimodal baseline model. As part of future work, we plan to perform further validations of our proposed model to estimate its generalizability on more neuroimaging datasets and utilize the patient-level deviation maps to further investigate heterogeneity of other neurodegenerative and neuropsychiatric disorders.

6 Acknowledgement

The preparation of this report was supported by the Centene Corporation contract (P19-00559) for the Washington University-Centene ARCH Personalized Medicine Initiative and the National Institutes of

Health (NIH) (R01-AG067103). Computations were performed using the facilities of the Washington University Research Computing and Informatics Facility, which were partially funded by NIH grants S10OD025200, 1S10RR022984-01A1 and 1S10OD018091-01. Additional support is provided The McDonnell Center for Systems Neuroscience.

Data used in preparation of this article were obtained from the Alzheimer's disease Neuroimaging Initiative (ADNI) database (http://adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found here

References

[1] Ana Lawry Aguila, Alejandra Jayme, Nina Montaña-Brown, Vincent Heuveline, and Andre Altmann. Multi-view-ae: A python package for multi-view autoencoder models. *Journal of Open Source Software*, 8(85):5093–5093, 2023.

[2] Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, 57(1):289–300, 1995.

[3] Yanshuai Cao and David J Fleet. Generalized product of experts for automatic and principled fusion of gaussian process predictions. *arXiv preprint arXiv:1410.7827*, 2014.

[4] Aoyan Dong, Jon B Toledo, Nicolas Honnorat, Jimit Doshi, Erdem Varol, Aristeidis Sotiras, David Wolk, John Q Trojanowski, Christos Davatzikos, and Alzheimer's Disease Neuroimaging Initiative. Heterogeneity of neuroanatomical patterns in prodromal alzheimer's disease: links to cognition, progression and biomarkers. *Brain*, 140(3):735–747, 2017.

[5] Bruce Fischl, David H Salat, Evelina Busa, Marilyn Albert, Megan Dieterich, Christian Haselgrove, Andre Van Der Kouwe, Ron Killiany, David Kennedy, Shuna Klaveness, et al. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron*, 33(3):341–355, 2002.

[6] Sayantan Kumar, Philip RO Payne, and Aristeidis Sotiras. Normative modeling using multimodal variational autoencoders to identify abnormal brain volume deviations in alzheimer's disease. In *Medical Imaging 2023: Computer-Aided Diagnosis*, volume 12465, page 1246503. SPIE, 2023.

[7] Andre F Marquand, Seyed Mostafa Kia, Mariam Zabihi, Thomas Wolfers, Jan K Buitelaar, and Christian F Beckmann. Conceptualizing mental disorders as deviations from normative functioning. *Molecular psychiatry*, 24(10):1415–1424, 2019.

[8] Andre F Marquand, Iead Rezek, Jan Buitelaar, and Christian F Beckmann. Understanding heterogeneity in clinical cohorts using normative models: beyond case-control studies. *Biological psychiatry*, 80(7):552–561, 2016.

[9] Sona Matloubi, Ali Mohammadzadeh, Zahra Jafari, and Alireza Akbarzade Baghban. Effect of background music on auditory-verbal memory performance. *Audiology*, 23(5), 2014.

[10] Richard C Mohs, David Knopman, Ronald C Petersen, Steven H Ferris, Chris Ernesto, Michael Grundman, Mary Sano, Linas Bieliauskas, David Geldmacher, Chris Clark, et al. Development of cognitive instruments for use in clinical trials of antidementia drugs: additions to the alzheimer's disease assessment scale that broaden its scope. *Alzheimer disease and associated disorders*, 1997.

[11] Susanne G Mueller, Michael W Weiner, Leon J Thal, Ronald C Petersen, Clifford R Jack, William Jagust, John Q Trojanowski, Arthur W Toga, and Laurel Beckett. Ways toward an early diagnosis in alzheimer's disease: the alzheimer's disease neuroimaging initiative (adni). *Alzheimer's & Dementia*, 1(1):55–66, 2005.

[12] Walter HL Pinaya, Cristina Scarpazza, Rafael Garcia-Dias, Sandra Vieira, Lea Baecker, Pedro F da Costa, Alberto Redolfi, Giovanni B Frisoni, Michela Pievani, Vince D Calhoun, et al. Using normative modelling to detect disease progression in mild cognitive impairment and alzheimer's disease in a cross-sectional multi-cohort study. *Scientific reports*, 11(1):1–13, 2021.

[13] Cathie Sudlow, John Gallacher, Naomi Allen, Valerie Beral, Paul Burton, John Danesh, Paul Downey, Paul Elliott, Jane Green, Martin Landray, et al. Uk biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS medicine*, 12(3):e1001779, 2015.

[14] Mike Wu and Noah Goodman. Multimodal generative models for scalable weakly-supervised learning. *Advances in Neural Information Processing Systems*, 31, 2018.