MODALITY-AGNOSTIC STYLE TRANSFER FOR HOLISTIC FEATURE IMPUTATION

Seunghun Baek[1*] *Jaeyoon Sim*[1*] *Mustafa Dere*[2] *Minjeong Kim*[3] *Guorong Wu*[2] *Won Hwa Kim*[1]

[1] Pohang University of Science and Technology, Pohang, South Korea
[2]University of North Carolina at Chapel Hill, Chapel Hill, USA
[3] University of North Carolina at Greensboro, Greensboro, USA

ABSTRACT

Characterizing a preclinical stage of Alzheimer's Disease (AD) via single imaging is difficult as its early symptoms are quite subtle. Therefore, many neuroimaging studies are curated with various imaging modalities, e.g., MRI and PET, however, it is often challenging to acquire all of them from all subjects and missing data become inevitable. In this regards, in this paper, we propose a framework that generates unobserved imaging measures for specific subjects using their existing measures, thereby reducing the need for additional examinations. Our framework transfers modality-specific style while preserving AD-specific content. This is done by domain adversarial training that preserves modality-agnostic but AD-specific information, while a generative adversarial network adds an indistinguishable modality-specific style. Our proposed framework is evaluated on the Alzheimer's Disease Neuroimaging Initiative (ADNI) study and compared with other imputation methods in terms of generated data quality. Small average Cohen's $d < 0.19$ between our generated measures and real ones suggests that the synthetic data are practically usable regardless of their modality type.

1. INTRODUCTION

Early diagnosis of an irreversible neurodegenerative disease such as Alzheimer's Disease (AD) is critical to delay its progression. Images are often used to observe structures and functions of the brain, however, the diagnosis with a single imaging modality [1, 2] is challenging as the changes in preclinical stages such as Mild Cognitive Impairment (MCI) are very subtle. Using multiple modalities will definitely help improve confidence in the diagnosis, as the imaging scans are based on different underlying mechanisms and reflect different aspects of the disease such as brain structure (e.g., cortical thickness), metabolism (e.g., FDG) and protein accumulation (e.g., Tau and β-amyloid).

Despite the usefulness of adopting multiple modalities, it is practically infeasible as obtaining various imaging scans is time-consuming, costly, and burdening to each subject. This can be a significant barrier for patients who are potentially at risk of developing AD from MCI, as the high cost of imaging will discourage them from medical attention. In practice, affordable magnetic resonance imaging (MRI) is first taken and, expensive Positron Emission Tomography (PET) scans with various tracers are recommended where MCI-specific changes are differently characterized by accumulation of proteins. In such a scenario, if we can accurately impute measures of unobserved modalities from existing ones, e.g., generating expensive fluorodeoxyglucose (FDG) measures from inexpensive MRI, then it would help better analysis of a patient without going through burdening imaging protocols at a significantly reduced cost.

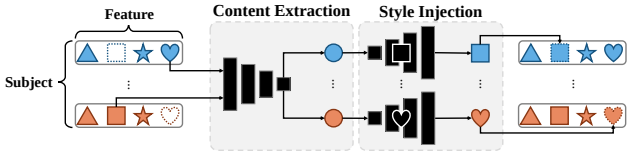

Fig. 1: An overview of our framework. In Content Extraction, modality-agnostic embedding is extracted from any type of feature for the subject. In Style Injection, modality-specific generators can generate missing features not present in the original subject. Shape: imaging scan (i.e., domain), Color: AD-stage label (i.e.,class).

Recent studies revealing high correlations between different imaging modalities for AD analyses demonstrate that the scenario above is potentially feasible [3, 4, 5, 6]. Several works [7, 8] focused on this correlation and tried to generate PET scans from MRI scans by directly applying basic generative adversarial network (GAN) or cycleGAN [9]. Other models such as conditional GAN (cGAN) [10] and Wasserstein GAN (WGAN) [11] can be an option, however, it is difficult to naively utilize them as the number of required generators will increase as a combination of the number of modalities. Traditional approaches such as Mean imputation [12] are still a golden standard even though they are unrealistic.

Although the progression of AD from MCI is differently characterized by different modalities and radiotracers, they should commonly contain AD-specific information at AD-specific regions of interest (ROI) which we consider as "content". We focus on this modality invariant content; we design an architecture that first projects a sample to a modality agnostic latent space, and the modality-agnostic embedding is put on with different realistic styles of various imaging modalities. With the proposed model, our goal is to universally generate various imaging measures for each subject that accurately reflects both the modality and AD-related information to impute missing data, without exhaustively training one-to-one mapping between different modalities.

*Equal contribution.

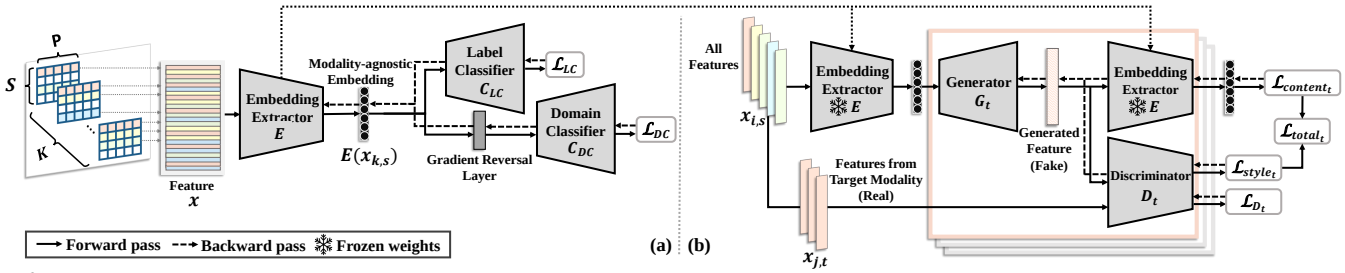

Fig. 2: Illustration of our framework. (a) Content extraction, (b) Style injection. In (a), our framework trains embedding extractor to obtain domain-agnostic embedding which still contain label-specific information. In (b), our framework endeavors to train generators from extracted domain-agnostic embedding to generate realistic measures.

Key Contributions: **1)** Our method generates probable estimation of unobserved imaging measures for specific subjects using their existing measures. **2)** To the best of our knowledge, our method is *the first* to address numerical data imputation via style transfer. **3)** Experimental results on ADNI data demonstrate that the data generated by our method can offer sufficiently realistic information for downstream analyses.

2. METHOD

Consider a dataset where each sample is expressed as multivariate measures taken from multiple imaging scans. If a subject skips some scans, which is common in many cohorts, then the sample ends up with missing data that discredits downstream analyses. To fully utilize such data, in this section, we propose a novel framework that completes the measures from missing modalities by imputing those for each subject based on their existing measurements. This is done by extracting modality-agnostic embedding with content only (e.g., target variable), and then training generators that equip realistic representations according to a specific modality. Our proposed framework consists of two consecutive parts: (i) modality-agnostic content extraction, and (ii) modality-wise style injection as illustrated in Fig. 1, and their details are given below.

2.1. Phase1 - Content Extraction

In this phase, our framework aims to train an embedding extractor E to extract a modality-agnostic embedding from data. This embedding has semantic component as a 'content' that is task-specific information (e.g., AD-specific biomarker) that is invariant to the domain-shift (e.g., modality). For this, domain adversarial training [13] is executed to make E learn modality-agnostic latent space.

Consider a subject going through S different imaging scans, and P-dimensional measurements (from P brain regions) are derived from individual scans. Then, the k-th sample is given as $x_k \in \mathbb{R}^{S\times P}$, where $x_{k,s} \in \mathbb{R}^P$ is denoted as a modality-specific data for the k-th subject and y_k is its corresponding diagnostic label. If the subject skips s-th scan, the whole $x_{k,s}$ becomes missing. To train E to extract modality-invariant disease-specific contents, individual $x_{k,s}$ is used to train the E to become agnostic on s.

The embedding $E(x_{k,s})$ is made modality-agnostic using AD-diagnostic label classifier C_{LC} and modality classifier C_{DC}. The $E(x_{k,s_i})$ and $E(x_{k,s_j})$ should be similar where $s_i \neq s_j$ and $s_i, s_j \in \{1, 2, \cdots, S\}$ for a specific subject k. Concurrently, C_{LC} is expected to accurately predict class labels, whereas C_{DC} should exhibit uncertainty in distinguishing the type of imaging scans. Thus, we use two loss functions $\mathcal{L}_{LC}$ and $\mathcal{L}_{DC}$ corresponding to C_{LC} and C_{DC} respectively:

$$\mathcal{L}_{LC} = \mathcal{J}(y_k, C_{LC}(E(x_{k,s}))), \;\; \mathcal{L}_{DC} = \mathcal{J}(s, C_{DC}(E(x_{k,s}))) \quad (1)$$

where $\mathcal{J}$ is a loss function (e.g. Cross-Entropy). While C_{LC} and C_{DC} are trained from $\mathcal{L}_{LC}$ and $\mathcal{L}_{DC}$ independently, E is updated in the direction of decreasing $\mathcal{L}_{LC}$ and increasing $\mathcal{L}_{DC}$. We adopt the gradient reversal layer [13] that converts the sign of gradient between E and C_{DC} to make E remove the modality-specific information, such that its imaging modality cannot be distinguished by C_{DC} but still accurately classifiable on its diagnostic label by C_{LC}.

2.2. Phase 2 - Style Injection

In the context of style transfer [14, 15], the term 'style' refers to modality-specific information that encompasses various aspects of data distinct from the content. In section 2.1, content was defined as task-specific information that is invariant to the modality type. Thus, style needs to contain a modality-specific information which is shared among data from the same modality. To combine the given content with a target style for the imputation, style needs to be added to the content without affecting the original content. Notice that the number of pairs increases as a combination of the number of styles, i.e., modalities, if we were to naively train pair-wise generation between different modalities. However, as shown in Fig. 2(b), because we extract modality-agnostic representation from 2.1, we only need to train S number of generator and discriminator pairs (G_t, D_t) where $t \in \{1, 2, \cdots S\}$ denotes a specific modality type.

Training a (G_t, D_t) pair requires separate inputs. To train G_t, all existing $x_{i,s}$, where $i \in \{1, 2, \cdots, K\}$ and $s \in \{1, 2, \cdots, S\}$, are fed to E to produce $E(x_{i,s})$, which is used as a seed to generate fake sample $G_t(E(x_{i,s}))$ which facilitates training of G_t. To train D_t, all real $x_{j,t}$ are taken conditioned on a target modality t, if j-th sample has t-th feature. The $x_{j,t}$ and $G_t(E(x_{i,s}))$ are fed into D_t for real-and-fake discrimination, and G_t is trained to deceive D_t by generating realistic data that mimics the target modality. At

the same time, D_t is trained adversarially to differentiate real and fake. This can be expressed using the minimax loss as

$$\min_{G_t}\max_{D_t}\mathbb{E}_{x\sim p_{\text{data}}(x_{j,t})}[\log D_t(x)]+\mathbb{E}_{z\sim p_{\text{data}}(x_{i,s})}[1-\log D_t(G_t(E(z)))]. \quad (2)$$

To optimize Eq. (2), each loss for training D_t and G_t are formulated as

$$\mathcal{L}_{D_t}(x_{j,t}, x_{i,s}) = -\log D_t(x_{j,t}) - \log\left(1 - D_t(G_t(E(x_{i,s})))\right), \quad (3)$$

$$\mathcal{L}_{G_t}(x_{i,s}) = -\log D_t(G_t(E(x_{i,s}))). \quad (4)$$

To ensure that the generated measures by G_t retain the content with a target modality style, $\mathcal{L}_{content}$ is defined as a $L2$-norm between the original embedding $E(x_{i,s})$ and that of the generated data $\hat{x}_{i,t} = G_t(E(x_{i,s}))$ given by

$$\mathcal{L}_{content} = \|E(x_{i,s}) - E(\hat{x}_{i,t})\|_2. \quad (5)$$

By minimizing Eq (5), G_t is guided to preserve content while generating a sample of a particular modality. Our full objective for training G_t is a combination of $\mathcal{L}_{G_t}$ and $\mathcal{L}_{content}$ as

$$\mathcal{L}_{total} = \alpha\mathcal{L}_{G_t} + \beta\mathcal{L}_{content} \quad (6)$$

where α and β are user-parameters. With $\mathcal{L}_{total}$, each generator is enabled to incorporate modality-specific style into the given content.

2.3. Imputation Procedure

Assume that a subject k has missing feature x_{k,s_i}. Our framework can generate $\hat{x}_{k,s_i}$ from an existing feature x_{k,s_j} for $s_i \neq s_j$, utilizing the trained content extractor E and style generator G_{s_i}. In this way, we ensure that all K subjects have every feature for S modalities for full utilization of the data.

3. EXPERIMENTAL RESULTS

3.1. Datasets

We evaluate our method on four imaging measures from MRI and PET scans from the ADNI study [16]. Each image was partitioned into 148 cortical and 12 sub-cortical regions using Destrieux atlas [17]. For each parcellation, region-specific imaging features including Standard Uptake Value Ratio (SUVR) [18] of metabolic activity from FDG-PET, β-amyloid protein from Amyloid-PET (β-Amy), Tau protein from Tau-PET and cortical thickness (CT) from MRI were measured. Cerebellum [19] was used as the reference region to standardize the SUVR values across the brain. Table 1 presents the demographics for each measure, including subjects who have taken all the imaging scans. A total of N=222 subjects have all scanned data from MRI to PET, which served as the baseline. The diagnostic labels for each subject include cognitive normal (CN), early mild cognitive impairment (EMCI), and late mild cognitive impairment (LMCI), which we used to design the 3-way classification.

Table 1: Sample-size of ADNI dataset for Different Modalities.

Label	Modality				
	CT	Tau	FDG	β-Amy	Common Subjects
CN	805	237	861	735	97
EMCI	486	186	597	833	87
LMCI	248	105	1138	447	38

3.2. Experiment setup

Baselines. We utilized **Mean** imputation [12] as a classical method, which imputed the missing values by replacing them with the class-specific means for each imaging scan. Prominent generative models such as **cGAN** [10] and **WGAN** [11] also serve as baselines. To apply the generative models for imputation, we trained generators for all pairs (i.e., 16 pairs) of imaging modalities.

Training. For E, C_{LC}, C_{DC}, G_t and D_t, Multi-Layer Perceptron (MLP) with 5, 2, 2, 13 and 5 layers were chosen. We used AdamW [20] optimizer with learning rate 10^{-3}, except for D_t with 10^{-5}. Weight decay at 0.01 was adopted for every linear layer. In Phase 1, we set the embedding size to 256 and trained E, C_{LC} and C_{DC} with 8×10^3 epoches. In Phase 2, each G_t and D_t pair was trained with 3×10^4 epoches in a 9-to-1 ratio iteratively. In L_{total}, α and β were 1 and 100.

Evaluation. We utilized MLP with 2, 3 and 4 layers as our backbone network, progressively getting larger, for the downstream classification denoted as 2-MLP, 3-MLP and 4-MLP respectively. For training of phase 1 and 2 (for E, G_t and D_t), only the samples with missing measures were used to avoid double dipping for a downstream task. In the downstream classification, samples with missing data using imputation methods including ours were utilized during the training. The experiments were performed with 5-fold cross validation (CV) to obtain unbiased results.

To evaluate the performance, we employed multi-class accuracy, weighted precision and recall averaged across the CV. Each baseline was implemented and trained to achieve optimal outcomes for fair comparisons. To verify whether the generated data are realistic, we computed the similarity between the actual and generated measurements for each modality using Cohen's d [21].

3.3. Quantitative Results

The averaged Cohen's d over all ROIs are reported in Table 2, and their visualization on EMCI subjects are given in Fig. 3. The effect sizes of the data generated by our modality-specific

Table 2: The average absolute Cohen's d across all the ROIs between actual and generated distributions. Lower values are better, and the values ≤ 0.2 are in **bold**.

Modality (Actual)	Modality (Generated)	Ours			cGAN [10]			WGAN [11]		
		CN	EMCI	LMCI	CN	EMCI	LMCI	CN	EMCI	LMCI
CT	CT	0.238	0.219	0.239	0.458	0.516	0.496	0.414	**0.192**	**0.162**
	Tau	**0.122**	**0.142**	0.358	0.315	0.425	0.401	0.212	**0.081**	0.222
	FDG	**0.141**	0.216	0.231	0.216	0.380	0.431	0.255	0.274	**0.063**
	β-Amy	**0.105**	**0.135**	0.211	0.400	0.296	0.460	**0.108**	**0.105**	0.248
Tau	CT	**0.192**	0.245	0.259	0.520	0.466	0.471	**0.085**	0.237	0.432
	Tau	**0.138**	**0.129**	0.379	**0.151**	**0.149**	0.224	0.349	**0.116**	0.369
	FDG	**0.133**	0.206	**0.185**	0.346	0.378	0.635	**0.103**	**0.132**	**0.158**
	β-Amy	**0.088**	**0.107**	**0.121**	**0.194**	0.309	**0.173**	**0.123**	0.271	0.461
FDG	CT	0.237	0.217	**0.167**	0.572	0.644	0.512	**0.168**	**0.188**	**0.193**
	Tau	**0.158**	**0.140**	0.447	0.373	0.423	0.371	**0.161**	0.458	0.487
	FDG	**0.142**	0.213	**0.146**	0.442	0.442	0.297	0.228	0.209	0.250
	β-Amy	**0.110**	**0.106**	**0.167**	0.265	0.300	0.415	**0.161**	**0.150**	0.411
β-Amy	CT	0.209	**0.185**	0.288	0.604	0.665	0.872	**0.186**	0.486	0.716
	Tau	**0.139**	**0.139**	0.391	0.267	0.333	0.290	0.423	0.645	0.888
	FDG	**0.142**	**0.189**	**0.178**	0.368	0.496	0.629	**0.076**	**0.166**	**0.135**
	β-Amy	**0.152**	**0.103**	**0.137**	0.372	0.386	0.402	0.249	**0.063**	0.292
Mean		**0.152**	**0.168**	0.244	0.366	0.413	0.442	0.206	0.235	0.342

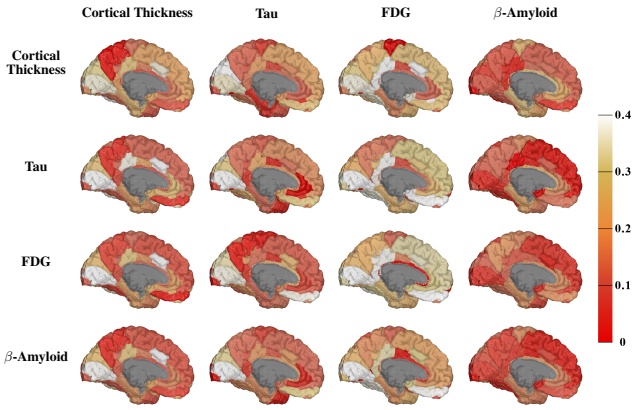

Fig. 3: Visualization of the averaged absolute Cohen's d between actual distribution and generated distribution on the inner left cortical regions of EMCI subjects. AD-specific regions show better imputation results as lower Cohen's d implies higher correspondence. (Row: Source, Column: Target.)

generators were small, i.e., 0.188 on average, far less than 0.407 and 0.261 from individually trained cGAN and WGAN, demonstrating the feasibility of our generated data.

The primary difference between our method and other generative models is the number of generators required for imputation. Unlike our framework only needs S generators through modality-agnostic embedding, other conventional generative methods need S^2 generators for the same task. When the sample size is small, training multiple generators can be challenging, whereas our method successfully trains the generators using modality-agnostic embeddings.

3.4. Qualitative Results

In Fig. 4, we visualized a generated sample (standardized at each ROI), along with the true measurements from a subject in the CN group. We selected a subject who underwent all imaging scans and was not used during generator training to provide ground truth for our estimation. From CT of this subject, CT, Tau, FDG and β-amyloid measures were generated and compared with the ground truth. Although we did not train the model with modality pairs from the same subject as they are insufficient, our model generated every ROI accurately toward observed measurement, as visualized in Fig. 4 (bottom).

3.5. Downstream MCI Classification Performance

In Table 3, we reported the performance of the downstream classification with 5-fold CV. The MLP using imputed data from our framework outperformed all other baselines in accuracy and weighted recall across every MLP classifiers. Our approach showed an accuracy increase over the baseline in 2-MLP (0.022), 3-MLP (0.026) and 4-MLP (0.040), indicating that generated samples enable to train larger models. However, as we stacked more layers for classifier, there was little increase in performance over the baseline with mean imputation, indicating that its generated data are limited. While usability of mean imputation was limited to

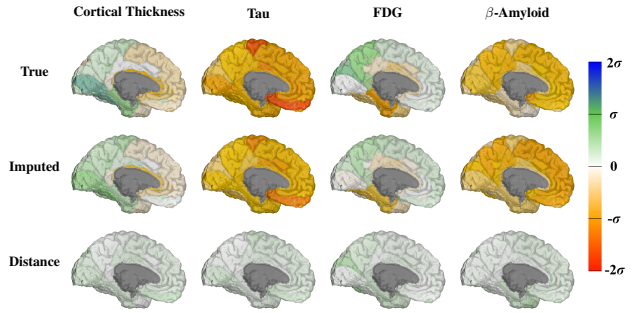

Fig. 4: Visual comparison of observed measurement (Top) and our estimation (Middle) on the inner view of left hemisphere from a CN subject. Our estimations were generated from observed CT measurement of the subject. All measurements were standardized, and the distance between ground truth and generated result are given at the bottom. BrainPainter [22] was used to generate the drawings.

downstream classification, our primary competitors were cGAN and WGAN as they generate subject-wise estimation. Regardless of the model size, our model showed superior performance in terms of accuracy and weighted recall over cGAN and WGAN. This could be because the two generative baselines were trained using fewer samples and lacked any precautions against mode-collapse, unlike our approach that incorporates content loss.

Table 3: Classification performance (CN/EMCI/LMCI) on ADNI data with all modalities. The number of generators required are given in ().

Model	Method	Imputation Concept	Performance		
			Accuracy	Precision	Recall
2-MLP	Baseline	No Imputation	0.776 ± 0.032	0.794 ± 0.030	0.776 ± 0.032
	Mean	Representative Value	0.792 ± 0.081	0.806 ± 0.052	0.792 ± 0.081
	cGAN [10]	Generative Model (16)	0.784 ± 0.092	0.743 ± 0.183	0.784 ± 0.092
	WGAN [11]	Generative Model (16)	0.797 ± 0.075	**0.834** ± 0.072	0.797 ± 0.075
	Ours	Generative Model (4)	**0.798** ± 0.040	0.789 ± 0.048	**0.798** ± 0.040
3-MLP	Baseline	No Imputation	0.785 ± 0.039	**0.811** ± 0.032	0.785 ± 0.039
	Mean	Representative Value	0.802 ± 0.096	0.790 ± 0.087	0.802 ± 0.096
	cGAN [10]	Generative Model (16)	0.805 ± 0.076	0.809 ± 0.095	0.805 ± 0.076
	WGAN [11]	Generative Model (16)	0.784 ± 0.072	0.770 ± 0.093	0.784 ± 0.072
	Ours	Generative Model (4)	**0.811** ± 0.061	0.791 ± 0.076	**0.811** ± 0.061
4-MLP	Baseline	No Imputation	0.785 ± 0.018	**0.804** ± 0.018	0.785 ± 0.018
	Mean	Representative Value	0.788 ± 0.084	0.771 ± 0.092	0.788 ± 0.084
	cGAN [10]	Generative Model (16)	0.788 ± 0.105	0.766 ± 0.105	0.788 ± 0.105
	WGAN [11]	Generative Model (16)	0.792 ± 0.054	0.796 ± 0.093	0.792 ± 0.056
	Ours	Generative Model (4)	**0.825** ± 0.018	0.793 ± 0.051	**0.825** ± 0.018

4. CONCLUSION

In this paper, we propose a novel framework that generates unobserved imaging measures for specific subjects using their existing measures. To reduce the need for taking several imaging scans, our framework addresses the imputation of missing measures by transferring modality-specific style while preserving AD-specific content. Experimental results on the ADNI study show that our model provides a probable estimation of target modality for individual subjects, which yields similar distributions of generated measurements to those from observed data and helps downstream analyses. Since our work is applicable regardless of modality type, our approach has the potential to be adopted by other neuroimaging studies that are limited by missing measures.

5. ACKNOWLEDGEMENT

This research was supported by NRF-2022R1A2C2092336 (50%), IITP-2022-0-00290 (20%), IITP-2019-0-01906 (AI Graduate Program at POSTECH, 10%) funded by MSIT, HU22C0171 (10%) and HU22C0168 (10%) funded by MOHW from South Korea.

6. REFERENCES

[1] Seunghun Baek, Injun Choi, et al., "Learning covariance-based multi-scale representation of neuroimaging measures for alzheimer classification," in *ISBI*. IEEE, 2023, pp. 1–5.

[2] Jaeyoon Sim, Sooyeon Jeon, InJun Choi, Guorong Wu, and Won Hwa Kim, "Learning to approximate adaptive kernel convolution on graphs," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024, vol. 38, pp. 4882–4890.

[3] Rik Ossenkoppele, Smith, et al., "Associations between tau, aβ, and cortical thickness with cognition in alzheimer disease," *Neurology*, vol. 92, no. 6, pp. e601–e612, 2019.

[4] Theresa M Harrison, Du, et al., "Distinct effects of beta-amyloid and tau on cortical thickness in cognitively healthy older adults," *Alzheimer's & dementia*, vol. 17, no. 7, pp. 1085–1096, 2021.

[5] Anna Rubinski, Nicolai Franzmeier, Julia Neitzel, Michael Ewers, and Alzheimer's Disease Neuroimaging Initiative (ADNI), "Fdg-pet hypermetabolism is associated with higher tau-pet in mild cognitive impairment at low amyloid-pet levels," *Alzheimer's Research & Therapy*, vol. 12, pp. 1–12, 2020.

[6] Jennifer L Whitwell, Graff-Radford, et al., "Imaging correlations of tau, amyloid, metabolism, and atrophy in typical and atypical alzheimer's disease," *Alzheimer's & Dementia*, vol. 14, no. 8, pp. 1005–1014, 2018.

[7] Jin Zhang, Xiaohai He, Linbo Qing, Feng Gao, and Bin Wang, "Bpgan: Brain pet synthesis from mri using generative adversarial network for multi-modal alzheimer's disease diagnosis," *Computer Methods and Programs in Biomedicine*, vol. 217, pp. 106676, 2022.

[8] Yongsheng Pan, Mingxia Liu, Chunfeng Lian, Tao Zhou, Yong Xia, and Dinggang Shen, "Synthesizing missing pet from mri with cycle-consistent generative adversarial networks for alzheimer's disease diagnosis," in *MICCAI*. Springer, 2018, pp. 455–463.

[9] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *ICCV*, 2017, pp. 2223–2232.

[10] Mehdi Mirza and Simon Osindero, "Conditional generative adversarial nets," 2014.

[11] Martin Arjovsky, Soumith Chintala, and Léon Bottou, "Wasserstein generative adversarial networks," in *International conference on machine learning*. PMLR, 2017, pp. 214–223.

[12] A Rogier T Donders, Geert JMG Van Der Heijden, Theo Stijnen, and Karel GM Moons, "A gentle introduction to imputation of missing values," *Journal of clinical epidemiology*, vol. 59, no. 10, pp. 1087–1091, 2006.

[13] Yaroslav Ganin, Ustinova, et al., "Domain-adversarial training of neural networks," *JMLR*, vol. 17, no. 1, pp. 2096–2030, 2016.

[14] Leon A Gatys, Alexander S Ecker, and Matthias Bethge, "Image style transfer using convolutional neural networks," in *CVPR*, 2016, pp. 2414–2423.

[15] Fujun Luan, Sylvain Paris, Eli Shechtman, and Kavita Bala, "Deep photo style transfer," in *CVPR*, 2017, pp. 4990–4998.

[16] Clifford R Jack Jr, Bernstein, et al., "The alzheimer's disease neuroimaging initiative (adni): Mri methods," *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 27, no. 4, pp. 685–691, 2008.

[17] Christophe Destrieux, Bruce Fischl, Anders Dale, and Eric Halgren, "Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature," *Neuroimage*, vol. 53, no. 1, pp. 1–15, 2010.

[18] Joseph A Thie, "Understanding the standardized uptake value, its methods, and implications for usage," *Journal of Nuclear Medicine*, vol. 45, no. 9, pp. 1431–1434, 2004.

[19] Mark Rapoport, Robert van Reekum, and Helen Mayberg, "The role of the cerebellum in cognition and behavior: a selective review," *The Journal of neuropsychiatry and clinical neurosciences*, vol. 12, no. 2, pp. 193–198, 2000.

[20] Ilya Loshchilov and Frank Hutter, "Decoupled weight decay regularization," 2017.

[21] Marnie E Rice and Grant T Harris, "Comparing effect sizes in follow-up studies: Roc area, cohen's d, and r," *Law and human behavior*, vol. 29, pp. 615–620, 2005.

[22] Răzvan V Marinescu et al., "Brainpainter: A software for the visualisation of brain structures, biomarkers and associated pathological processes," in *Multimodal Brain Image Analysis and Mathematical Foundations of Computational Anatomy: 4th International Workshop, MBIA*. Springer, 2019, pp. 112–120.