Acquiring a low-dimensional, environment-independent representation of brain MR images for content-based image retrieval

Shuya Tobari[1], Kenichi Oishi[2], Hitoshi Iyatomi[1]
for the Alzheimer's Disease Neuroimaging Initiative* and the Parkinson's Progression Markers Initiative.
[1]*Department of Applied Informatics, Graduate School of Science and Engineering, Hosei University, Tokyo, Japan*
[2]*Department of Radiology and Radiological Science, Johns Hopkins University School of Medicine, Baltimore, USA*
shuya.tobari.7b@gmail.com, koishi2@jhmi.edu, iyatomi@hosei.ac.jp

***Abstract*—To make content-based image retrieval (CBIR) technology for magnetic resonance (MR) images of the brain practical and useful for diagnosis and research, it is important to obtain low-dimensional representations that embody pathological attributes. However, recent evidence suggests that variations in domains resulting from differences in imaging equipment and protocols at each imaging facility can overshadow pathological attributes. In this study, we propose a novel approach known as multidecoder adversarial domain adaptation (MD-ADA) to obtain low-dimensional representations of brain MR images that preserve pathological features while mitigating domain differences. This method combines adversarial domain adaptation techniques with convolutional autoencoders that have distinct decoders for each domain and employs adversarial learning to prevent domain discrimination from the produced low-dimensional representations. Experimental evaluations on two datasets, ADNI and PPMI, comprising 4,168 brain images demonstrate that the proposed MD-ADA significantly reduces domain differences between datasets without compromising the recoverability of brain images or the accuracy of disease classification.**

***Index Terms*—ADNI, PPMI, CBIR, domain harmonization, dimensional reduction, 3D brain MRI**

I. Introduction

Magnetic resonance (MR) images and their associated clinical information are stored in picture archiving and communication systems (PACS) [1] to facilitate the centralized management of scan images. These images are then searched for diagnostic or research purposes, and querying or registering images in these databases commonly involves using keywords that represent the structural and clinical features of the brain. However, selecting appropriate keywords requires knowledge and experience, making it costly. Thus, content-based image retrieval (CBIR) that can search for images themselves is advantageous in medical settings [2].

From a machine learning perspective, brain MR images are high-dimensional data with a large number of voxels, and contain a lot of information, other than disease and biological aspects, that is noise to CBIR. Therefore, appropriate dimensionality reduction reflecting disease features is essential. Arai et al. [3] successfully obtained low-dimensional representations of high-resolution brain MR images, which are essential for achieving CBIR, using 3D convolutional autoencoders (3D-CAE). However, it has been confirmed that domain differences exist due to differences in scanners and protocols at each imaging site, which affect feature extraction based on biological features [4], [5]. Non-biological variations—such as magnetic field strength, scanner manufacturer, and reconstruction protocol at each facility—have different effects on the images and have a significant impact on subsequent processes such as classification and clustering [4]–[7]

To make CBIR based on a diverse range of brain MR images accumulated through a multicenter and long-term study, an extractor capable of robust feature extraction that can handle domain differences is required. In addition to statistical and image processing-based methods [6], [8]–[11] that have been proposed to address this problem, methods based on deep learning [4], [12]–[14] have recently been introduced.

Classical harmonization methods have been successful in harmonizing images by correcting the brightness distribution and histogram equalization in sub-regions of an image [8], but these methods can only approximate the statistics that can be computed from the images; these are not suitable for images with pathological features that affect the intensity profiles [4]. Methods based on ComBat [9], an empirical Bayesian method developed to eliminate batch effects in genetics, have been reported to successfully remove non-biological variation while preserving biological features [6], [10], [11]. ComBat models image features by incorporating both biological variables and scanner effects using multivariate linear mixed effects regression. Furthermore, these models have shown superior performance even with small sample sizes due to the use of empirical Bayesian methods to learn the model parameters. However, these models have limitations, such as (i) insufficient performance in complex mappings across multiple regions, (ii) assumptions of specific prior probabilities (Gaussian or inverse gamma) may not be appropriate, and (iii) susceptibility to outliers [15].

*Data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database(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 at https://adni.loni.usc.edu/wpcontent/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf

Deep learning-based methods for harmonizing brain MR images can be broadly classified into two approaches based on: 1) image generation [4], [12], [13] and 2) adversarial domain adaptation techniques [14].

The former is a domain harmonization method that uses generative models such as U-net [16] and generative adversarial networks (GANs) [17] to perform style transfer of images from any domain to a specific domain, thereby generating high-quality transformed images with excellent results reported. On the other hand, due to the constraint that the size and shape of the target structure (i.e., the brain in the image) cannot be significantly changed, it is not suitable for cases in which the size and shape of the brain differ between domains. In addition, because training the model is difficult and requires a large amount of data, many models are limited to implementation on 2D slices. Therefore, the possibility of missing context between adjacent slices cannot be eliminated.

The latter method is a means to get feature representations that is invariant to the domain of the data and can appropriately perform a given main task using adversarial domain adaptation techniques formalized by Ganin et al. [18]. Dinsdale et al. [14] applied this technique to the domain harmonization of MR images and demonstrated robust discrimination against the influence of domain differences in deep learning models that perform age prediction and segmentation. However, this method is limited to obtaining feature representations that retain partial information regarding the brain for a specific task such as classification, and is not designed to obtain low-dimensional feature representations that retain information regarding the entire brain image. Therefore, these methods are not suitable for CBIR, which requires appropriate dimensionality reduction that preserves information regarding the entire brain.

In this paper, we propose a new effective method, multidecoder adversarial domain adaptation (MD-ADA), to address the domain shift problem caused by differences in scanning equipment, protocols, and other environmental factors in imaging centers, which hinders feature extraction based on essential biological features necessary for CBIR implementation. MD-ADA is based on the method proposed by Dinsdale et al. [14] and applies adversarial domain adaptation to a 3D-CAE. The new key feature of MD-ADA is a domain-specific decoder that improves domain harmonization and stabilizes training to achieve the desired low-dimensional representation.

II. Related Work

Furthering Arai et al.'s work on dimensionality reduction of brain MR images [3], Onga et al. [5] introduced distance learning and improved the feature representation to be more coherent for different pathological conditions. Nishimaki et al [19] further improved the interpretability of features and localized brain regions using variational autoencoders [20]. However, these studies did not take into account differences in imaging sites, thus leaving room for improvement for practical CBIR with multicenter data.

One domain harmonization method based on image generation is the use of generative models such as GANs [17] and CycleGANs [21], as proposed in previous studies [4], [12], [13]. Liu et al [13] used GAN for style transformation and successfully generated superior harmonic MR images by mapping MR images collected from different sites into the same latent space and then inverse transforming them. In addition, Arai et al. [4] applied CycleGAN to convert 3D brain MR images acquired from different MRI scanners to pseudo-standard scanner images, followed by dimensionality reduction using 3D-CAE. The acquired low-dimensional representation revealed reduced domain bias and improved disease classification accuracy.

On the other hand, a recent technique called adversarial domain adaptation aims to optimize the training model for both the main task, such as classification, and domain classification; this causes domain identification to become impossible while maintaining accuracy in the main task. Ganin et al. [18] revealed that this behavior can be achieved by simply adding a domain classifier consisting of standard convolutional layers and gradient reversal layers to a feedforward neural network. Another adversarial domain adaptation method was proposed by Tzeng et al. [22]. This method alternates between training the optimal domain classifier for a given feature representation and minimizing a confusion loss to optimize the feature extractor to bring the softmax output of the domain classifier closer to a uniform distribution, thereby successfully confusing the domain classifier to the maximum extent. Dinsdale et al. [14] achieved domain harmonization of MR images for a specific task using this adversarial domain adaptation. They also demonstrated that this method is more suitable for domain harmonization tasks than other adversarial domain adaptation methods, as it more effectively acquires feature representations with evenly distributed information across domains. However, as mentioned above, their method cannot acquire a low-dimensional representation that preserves the input information. When their method was modified for CBIR realization so that all the input information was converted to a low-dimensional representation (e.g., by eliminating the U-Net bypass), there was still room for improvement in harmonic performance when the domain differences were large.

III. Multidecoder adversarial domain adaptation(MD-ADA)

In the CBIR system considered in this paper, the input 3D brain MR image is dimensionally compressed by a feature extractor, and the similarity calculation is performed in the low-dimensional feature space, where the images in the database have been pre-compressed by a same method. This paper proposes a new feature extractor, multidecoder adversarial domain adaptation (MD-ADA), which solves the key challenge of harmonizing domain differences when implementing CBIR. The key feature of our proposal is that it has a decoder for each harmonized domain, as described in detail below.

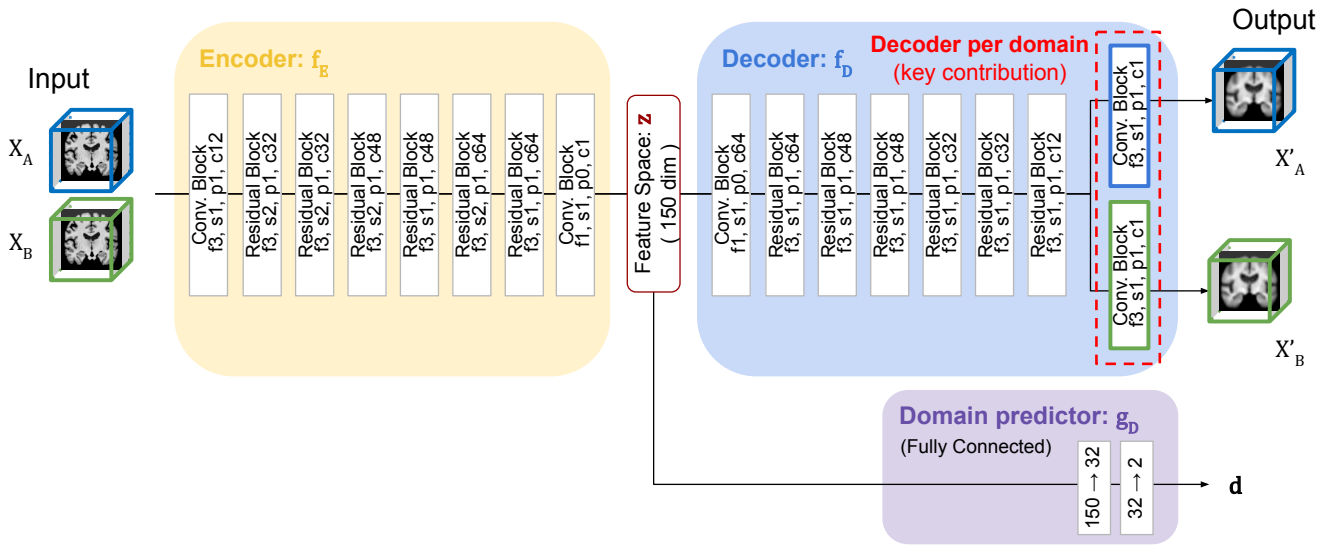

Fig. 1. Architecture of MD-ADA (two domain; f: filter size, s: stride size, p: padding size, c: output channel size).

A. Network architecture of MD-ADA

Figure 1 shows an example of a model that harmonizes two domains $\mathcal{A}$ and $\mathcal{B}$ using MD-ADA. MD-ADA is a machine learning model that builds on the low-dimensional compression of brain images, demonstrated in previous studies [3] [4], by incorporating insights from current machine learning techniques to improve the 3D-CAE encoder-decoder architecture. The MD-ADA model has a τ-shaped structure, similar to the approach of Dinsdale et al. [14], and includes a domain predictor g_D to perform adversarial domain adaptation. The Conv Block comprises convolution, batch normalization, and rectified linear unit (ReLU) activation. The Residual Block has a configuration published in the literature [34], combining multiple convolutions with skip connections to stabilize training as the number of layers increases. In addition, the Residual Block of the decoder performs dimension expansion using upsampling prior to application.

The encoder f_E transforms the input $\boldsymbol{x}$ of domain k ($\boldsymbol{x} \in X_k$) into a low-dimensional representation $\boldsymbol{z}$ for use in CBIR, while the decoder f_D aims to reconstruct an image $\boldsymbol{x'} \in X'_k$ that is identical to the input image. The domain predictor g_D is a training model that predicts the domain k to which $\boldsymbol{x}$ belongs based on $\boldsymbol{z}$ through supervised training. MD-ADA aim to obtain a low-dimensional representation $\boldsymbol{z}$ of the input image $\boldsymbol{x}$ that can be correctly reconstructed by the decoder, and to eliminate information that contributes to domain prediction from $\boldsymbol{z}$. And this elimination of domain-relevant information is achieved by training the encoder to produce the classification results $\boldsymbol{d}$ are uniform across all domains, even with a highly trained domain predictor g_D.

MD-ADA proposes an improvement to prevent the loss of harmonization capacity due to conflicting training, discussed below, by employing a decoder for each target domain. In the implementation, the decoder for each domain can be easily implemented by only differencing the final Conv Block. Therefore, the increase in the number of parameters is rather limited.

B. Training procedure

The training process of MD-ADA consists of three stages:

(i) Training the encoder f_E and decoder f_D to acquire the low-dimensional representation $\boldsymbol{z}$ that minimizes the reconstruction error with the input image.

(ii) Training only the domain predictor g_D to enable domain classification from the feature representation $\boldsymbol{z}$.

(iii) Fixing the domain predictor g_D and training only the encoder f_E to make the domain classification impossible from the feature representation $\boldsymbol{z}$.

For each stage, we optimize a different loss function, thereby leading to three consecutive steps ((i)–(iii)) in each training batch. The elimination of domain information is achieved through a combination of training in stages (ii) and (iii). In stage (ii), domain predictor g_D is optimized to perform domain classification from the low-dimensional representation $\boldsymbol{z}$ obtained in stage (i). In stage (iii), only encoder f_E is optimized, while keeping g_D fixed, to prevent it from accurately classifying the domain. As long as domain information remains in the feature representation $\boldsymbol{z}$, g_D can be optimized to enable domain classification, so this adversarial two-stage training can be repeated to eventually omit almost all domain information from $\boldsymbol{z}$, but leaving reconstructable information (brain information).

To train (i), the mean squared error, commonly used in CAE training, was calculated for each domain, and the macro average was used as the reconstruction error. For the training of (ii), cross-entropy—a loss function commonly used in class classification—was used. For the training of (iii), the confusion loss [22] shown in equation (1) was used, where n is the number of images in the mini-batch, K is the total

number of domains considered and k is K's index, and $p_{x,k}$ is the softmax output of the domain classifier for domain k. This loss becomes smaller when $p_{x,k}$ is evenly distributed across domains and larger when the output is biased toward a particular domain.

$$\mathcal{L}_{\mathrm{conf}} = -\frac{1}{n}\sum_{x=1}^{n}\sum_{k=1}^{K}\frac{1}{K}\log(p_{x,k}). \tag{1}$$

If sufficient harmonization is achieved between different domains and there is no domain information in the representation $\boldsymbol{z}$, it becomes difficult to reconstruct an original image that contains domain features when there is only one decoder, as in previous studies. In the proposed MD-ADA, by preparing a decoder for each domain, it is possible to reconstruct an input image that preserves domain information from a representation $\boldsymbol{z}$ that has lost domain information, thereby avoiding the training conflict between stages (i) and (iii).

IV. Experiments

A. Datasets and Preprocessing

Table I shows the details of the two datasets of T1-weighted brain MR images used in this experiment. The Alzheimer's Disease Neuroimaging Initiative 2 (ADNI2)[1] consists of four disease categories of increasing severity in the following order: cognitively normal (CN), mild cognitive impairment (MCI), late mild cognitive impairment (LMCI), and Alzheimer's disease (AD). The Parkinson's Progression Markers Initiative (PPMI)[2] consists of Parkinson's disease cases and healthy individuals. The CN category of ADNI2 and the control category of PPMI are medically equivalent, and since there are no anatomical features associated with Parkinson's disease on T1-weighted MRI, all data from PPMI were considered medically healthy (CN) and used to evaluate domain harmonization.

Further, all images were skull-stripped and linearly aligned to the JHU-MNI space [23] using MRICloud [24] provided by the collaborative research institution, Johns Hopkins University. The size of the images after processing by MRICloud was 181×217×181 pixels, but they were downsampled to 80×96×80 pixels for efficient training. MR images often have different intensity distributions and contrasts depending on the protocol or operator used during the scan; thus, the intensity normalization was performed. Specifically, the standard deviation σ of the intensity was calculated for each case, and negative values and values greater than 4σ were replaced by 0 and 4σ, respectively. Thereafter, a linear transformation was performed to normalize the values to have a maximum of 1 and a minimum of 0.

B. Evaluation

To confirm the effectiveness of the feature representation $\boldsymbol{z}$ generated by MD-ADA for CBIR, we conducted the following four evaluations. In all experiments, the dimension of the feature representation $\boldsymbol{z}$ was set to 150. Their details and settings are described in the following subsection.

[1] https://adni.loni.usc.edu [2] https://www.ppmi-info.org

TABLE I
Number of cases per dataset

Dataset	Case label	Number	Number of Patients
ADNI2	CN	1,121	227
	AD	674	192
	LMCI	926	143
	MCI	995	240
PPMI	Control	114	75
	PD	338	149

(1) Preservation of original information (*Reconstruction*)
(2) Domain information residuality (*Domain F1*)
(3) Preservation of disease information (*Diag F1*)
(4) Relevance of distance in feature space ($D_{CN-Ctrl}$)

To evaluate the effectiveness of the proposed MD-ADA, we conducted comparative evaluations using the following models for the four evaluation items described above:

(A) A simple dimension compression by 3D-CAE (baseline)
(B) Applying ComBat [9] to the feature representation obtained in (A) (+ ComBat)
(C) Adding Gaussian noise with mean 0 and standard deviations of 0.1, 0.03, and 0.01 to the feature representation obtained in (A) (+ Noise)
(D) Combining adversarial domain adaptation with (A) (+ ADA)
(E) The proposed method (+ MD-ADA)

Furthermore, we conducted five-fold cross-validation that considered patient labels to ensure that the same patients were not present in the training and evaluation data.

1) Preservation of original information (Reconstruction): To evaluate how well the acquired feature representation $\boldsymbol{z}$ retains original brain information, we calculated the root mean square error (RMSE) between the original image $\boldsymbol{x}$ and the reconstructed image $\boldsymbol{x'}$ obtained from the decoder. We also performed visual evaluation of the reconstructed brain images.

2) Domain information residuality (Domain F1): To evaluate the domain information residuals of the low-dimensional representation $\boldsymbol{z}$, we investigated the domain classification performance of $\boldsymbol{z}$ using L2-regularized logistic regression (Logistic-L2). In addition, for ADA and MD-ADA, we also evaluated the domain classification performance obtained from the domain predictor g_D of the final epoch. We used the F1 score as the evaluation metric. To ensure an accurate evaluation, we randomly selected one case from each patient for evaluation (the same applies to the following). Ideally, if appropriate domain harmonization is achieved, this value should be close to 1 divided by the number of domains (i.e., 0.5 in this experimental setting).

3) Preservation of disease information (Diag F1): To evaluate the preservation of disease information in $\boldsymbol{z}$, we constructed an Alzheimer's disease diagnostic model using logistic regression with L2 regularization based on $\boldsymbol{z}$ as described above, and then evaluated the diagnostic performance using the F1 score. To demonstrate that the acquired low-dimensional representation retains biologically meaningful information, it

TABLE II
COMPARISON OF THE EFFECTIVENESS OF THE FEATURE REPRESENTATION FROM EACH METHOD FOR CBIR.

		(1) Reconstruction (RMSE) ↓	(2) Domain F1 ↓ Logistic-L2	g_D	(3) Diag F1 ↑	(4) $D_{\mathrm{CN-Ctrl}}$ ↓
A)	3D-CAE (baseline)	**0.0853**	0.942	0.954	0.762	1.030
B)	+ Combat [9]	-	**0.524**	-	0.730	0.971
C)	+ Noise (SD: 0.01)	-	0.953	-	0.753	1.031
	+ Noise (SD: 0.03)	-	0.918	-	0.737	1.009
	+ Noise (SD: 0.1)	-	0.820	-	0.637	1.010
D)	+ ADA	0.0868	0.795	0.584	0.753	0.861
E)	+ MD-ADA	0.0861	0.702	**0.558**	**0.775**	**0.847**

is desirable for this value to be high. We excluded MCI and LMCI, which are cognitive impairment conditions, from the training and evaluation and treated PD and control of the PPMI dataset as CN, as mentioned earlier.

4) Relevance of distance in feature space ($D_{CN-Ctrl}$): To quantitatively evaluate the effect of the proposed harmonization, after each harmonization process, we evaluated the extent to which data groups in the same medical category were separated from each other in the low-dimensional feature space due to domain differences. Specifically, we evaluated the Mahalanobis distance $D_{CN-Ctrl}$ between the data of ADNI2-CN and the data of PPMI, which are considered medically similar. Furthermore, we normalized the Mahalanobis distance between AD and CN in the ADNI2 dataset to 1. $D_{CN-Ctrl}$ is expected to be smaller due to domain harmonization.

V. RESULTS

Table II summarizes the evaluation results. The best results in each category are indicated in bold. An example of a brain image and its reconstructions from low-dimensional representations are shown in Fig. 2. Combat(B) [9] and the method of adding noise to the low-dimensional representation (C) are methods that modify the feature representation itself. Therefore, reconstructed images are not available, and the quality of the image reconstruction is not evaluated.

The proposed MD-ADA (E) significantly reduced the domain information residuality (*Domain F1*), while slightly improving the disease classification accuracy (*Diag F1*) compared to other methods. It was also observed that the difference in the low-dimensional representation due to different datasets of the same disease ($D_{CN-Ctrl}$) was reduced. Furthermore, as also indicated in Fig. 2, no disadvantage of having multiple decoders was observed for the proposed method in image reconstruction from the low-dimensional representation, $\boldsymbol{z}$.

ComBat(B) successfully reduced domain classification accuracy (*Logistic-L2*) to an ideal value close to 50%, but it also reduced disease diagnosis ability (*Diag F1*) and eliminated important biological features for CBIR. The noise injection method (C), aimed at reducing domain classification ability, indicated some effect on domain harmonization, but at the significant expense of disease diagnostic ability. It also did not affect the data distribution in the feature representation and, thus, did not reduce the distance between datasets for the same group of cases ($D_{CN-Ctrl}$).

VI. DISCUSSION

In CBIR, obtaining a low-dimensional representation for input brain images is essential; this is achieved solely by the encoder. In the conventional ADA(D) with only one decoder, if domain information is omitted from the low-dimensional representation $\boldsymbol{z}$ obtained by the encoder, image reconstruction becomes difficult. This model learns to reduce the reconstruction error of the CAE that comprises an encoder–decoder; consequently, it becomes difficult to sufficiently reduce domain information from $\boldsymbol{z}$. In the proposed MD-ADA(E), appropriate reconstructed images can be obtained even after sufficient elimination of domain information from the low-dimensional representation $\boldsymbol{z}$ of input images by incorporating decoders for each domain. This leads to further elimination of domain information from $\boldsymbol{z}$, as confirmed by the results. Therefore, domain-specific decoders significantly contribute to achieving a superior encoder for obtaining the low-dimensional representation $\boldsymbol{z}$ of input images, even in cases where the image features are significantly different between domains. In addition, the domain-specific decoders are implemented with minimal increase in parameter count, mainly differing in a simple block with only one additional convolutional layer. Therefore, it is desirable for preventing overfitting.

VII. CONCLUSION

In this paper, we proposed MD-ADA—a training model that can both harmonize domain differences and compress dimensionality while capturing the biological features of the brain—to make CBIR for brain MR images. MD-ADA is a τ-shaped structure model that combines a 3D-CAE architecture with adversarial domain adaptation and has a decoder for each domain to be harmonized. We demonstrated that MD-ADA can preserve biological information, such as brain structure and disease characteristics, while achieving excellent domain harmonization. In the future, we aim to establish a superior domain harmonization technique by validating the effectiveness of MD-ADA using larger and more diverse datasets.

ACKNOWLEDGMENT

This research was supported in part by the Ministry of Education, Science, Sports and Culture of Japan (JSPS KAKENHI), Grant-in-Aid for Scientific Research (C), 21K12656, 2021-2023.

The MRI data collection and sharing for this project was funded by the Alzheimer's Disease Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2-0012). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from

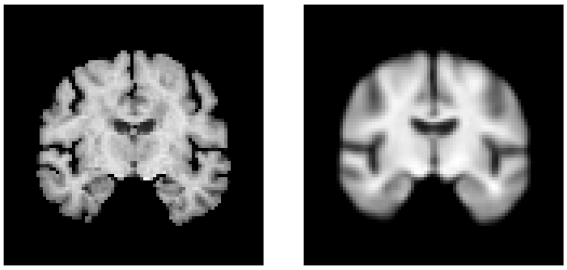

(a) Grand truth (b) Baseline (RMSE = 0.00863)

(c) ADA (RMSE = 0.00895) (d) MD-ADA (RMSE = 0.00883)

Fig. 2. Comparison of original image and reconstructed image in each method

the following: AbbVie, Alzheimer's Association; Alzheimer's Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; Bristol-Myers Squibb Company; CereSpir, Inc.; Cogstate; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.;Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Lumosity; Lundbeck; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer's Therapeutic Research Institute at the University of Southern California. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California. Additional MRI data used in the preparation of this article were obtained from the Parkinson's Progression Markers Initiative (PPMI) database (www.ppmi-info.org/data). For up-to-date information on the study, visit www.ppmi-info.org. PPMI—a public-private partnership—is funded by the Michael J. Fox Foundation for Parkinson's Research and funding partners, including AbbVie, Allergan, Avid Radiopharmaceuticals, Biogen, Biolegend, Bristol-Myers Squibb, Celgene, Denali, GE Healthcare, Genentech, GlaxoSmithKline, Lilly, Lundbeck, Merck, Meso Scale Discovery, Pfizer, Piramal, Prevail Therapeutics, Roche, Sanofi Genzyme, Servier, Takeda, Teva, UCB, Verily, Voyager Therapeutics, and Golub Capital.

REFERENCES

[1] R. H. Choplin, J. Boehme 2nd, and C. Maynard, "Picture archiving and communication systems: an overview." *Radiographics*, vol. 12, no. 1, pp. 127–129, 1992.

[2] A. Kumar, J. Kim, W. Cai, M. Fulham, and D. Feng, "Content-based medical image retrieval: a survey of applications to multidimensional and multimodality data," *Journal of digital imaging*, vol. 26, no. 6, pp. 1025–1039, 2013.

[3] H. Arai, Y. Chayama, H. Iyatomi, and K. Oishi, "Significant dimension reduction of 3D brain MRI using 3D convolutional autoencoders," in *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2018, pp. 5162–5165.

[4] H. Arai, Y. Onga, K. Ikuta, Y. Chayama, H. Iyatomi, and K. Oishi, "Disease-oriented image embedding with pseudo-scanner standardization for content-based image retrieval on 3D brain MRI," *IEEE Access*, vol. 9, pp. 165 326–165 340, 2021.

[5] Y. Onga, S. Fujiyama, H. Arai, Y. Chayama, H. Iyatomi, and K. Oishi, "Efficient feature embedding of 3D brain MRI images for content-based image retrieval with deep metric learning," in *2019 IEEE International Conference on Big Data (Big Data)*. IEEE, 2019, pp. 3764–3769.

[6] C. Wachinger, A. Rieckmann, S. Pölsterl, A. D. N. Initiative *et al.*, "Detect and correct bias in multi-site neuroimaging datasets," *Medical Image Analysis*, vol. 67, p. 101879, 2021.

[7] K. A. Clark, R. P. Woods, D. A. Rottenberg, A. W. Toga, and J. C. Mazziotta, "Impact of acquisition protocols and processing streams on tissue segmentation of T1 weighted MR images," *NeuroImage*, vol. 29, no. 1, pp. 185–202, 2006.

[8] Y. Gao, J. Pan, Y. Guo, J. Yu, J. Zhang, D. Geng, and Y. Wang, "Optimised MRI intensity standardisation based on multi-dimensional sub-regional point cloud registration," *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization*, vol. 7, no. 5-6, pp. 594–603, 2019.

[9] W. E. Johnson, C. Li, and A. Rabinovic, "Adjusting batch effects in microarray expression data using empirical Bayes methods," *Biostatistics*, vol. 8, no. 1, pp. 118–127, 2007.

[10] J.-P. Fortin, N. Cullen, Y. I. Sheline, W. D. Taylor, I. Aselcioglu, P. A. Cook *et al.*, "Harmonization of cortical thickness measurements across scanners and sites," *Neuroimage*, vol. 167, pp. 104–120, 2018.

[11] R. Pomponio, G. Erus, M. Habes, J. Doshi, D. Srinivasan, E. Mamourian *et al.*, "Harmonization of large MRI datasets for the analysis of brain imaging patterns throughout the lifespan," *NeuroImage*, vol. 208, p. 116450, 2020.

[12] B. E. Dewey, C. Zhao, J. C. Reinhold, A. Carass, K. C. Fitzgerald, E. S. Sotirchos *et al.*, "DeepHarmony: A deep learning approach to contrast harmonization across scanner changes," *Magnetic resonance imaging*, vol. 64, pp. 160–170, 2019.

[13] M. Liu, P. Maiti, S. Thomopoulos, A. Zhu, Y. Chai, H. Kim, and N. Jahanshad, "Style transfer using generative adversarial networks for multi-site mri harmonization," in *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part III 24*. Springer, 2021, pp. 313–322.

[14] N. K. Dinsdale, M. Jenkinson, and A. I. Namburete, "Deep learning-based unlearning of dataset bias for MRI harmonisation and confound removal," *NeuroImage*, vol. 228, p. 117689, 2021.

[15] F. Zhao, Z. Wu, L. Wang, W. Lin, S. Xia, D. Shen *et al.*, "Harmonization of infant cortical thickness using surface-to-surface cycle-consistent adversarial networks," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2019, pp. 475–483.

[16] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger, "3D U-Net: learning dense volumetric segmentation from sparse annotation," in *International conference on medical image computing and computer-assisted intervention*. Springer, 2016, pp. 424–432.

[17] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair *et al.*, "Generative adversarial networks," *Communications of the ACM*, vol. 63, no. 11, pp. 139–144, 2020.

[18] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette *et al.*, "Domain-adversarial training of neural networks," *The journal of machine learning research*, vol. 17, no. 1, pp. 2096–2030, 2016.

[19] K. Nishimaki, K. Ikuta, Y. Onga, H. Iyatomi, and K. Oishi, "Loc-VAE: Learning Structurally Localized Representation from 3D Brain MR Images for Content-Based Image Retrieval," in *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2022, pp. 2433–2438.

[20] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.

[21] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2223–2232.

[22] E. Tzeng, J. Hoffman, T. Darrell, and K. Saenko, "Simultaneous deep transfer across domains and tasks," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 4068–4076.

[23] K. Oishi, A. Faria, H. Jiang, X. Li, K. Akhter, J. Zhang *et al.*, "Atlas-based whole brain white matter analysis using large deformation diffeomorphic metric mapping: application to normal elderly and Alzheimer's disease participants," *Neuroimage*, vol. 46, no. 2, pp. 486–499, 2009.

[24] S. Mori, D. Wu, C. Ceritoglu, Y. Li, A. Kolasny, M. A. Vaillant *et al.*, "MRICloud: delivering high-throughput MRI neuroinformatics as cloud-based software as a service," *Computing in Science & Engineering*, vol. 18, no. 5, pp. 21–35, 2016.