Improved 3D Brain Generation with Cycle-Consistent Embedding GAN

Shibo Xing, Harsh Sinha, Seong Jae Hwang, University of Pittsburgh

Background and Task

Modern generative adversarial networks (GANs) have been enabling the realistic generationof full 3D brain images by sampling from a latent space prior Z(i.e., random vectors) and mapping it to realistic images in X(e.g., 3D MRIs). To address the ubiquitous mode collapse issue, recent works have strongly imposed certain characteristics such as Gaussianness to the prior by also explicitly mapping to encoder. These efforts, however, fail to accurately map 3D brain images to the desirable prior, which the generator assumes to besampling the random vectors from. On the other hand, Variational Auto-Encoding GAN(VAE-GAN) solves mode collapse by enforcing Gaussianness by two learned parameter,yet causes blurriness in images. In this work, we show how our cycle consistent embedding GAN (CCE-GAN) both accurately encodes 3D MRIs to the standard normal prior, and maintains the quality of the generated images. We achieve this without a network-based code discriminator via the Wasserstein measure. We quantitatively and qualitatively assessthe embeddings and the generated 3D MRIs using healthy T1-weighted MRIs from ADNI.

Variational Auto-Encoding (VAE) GAN

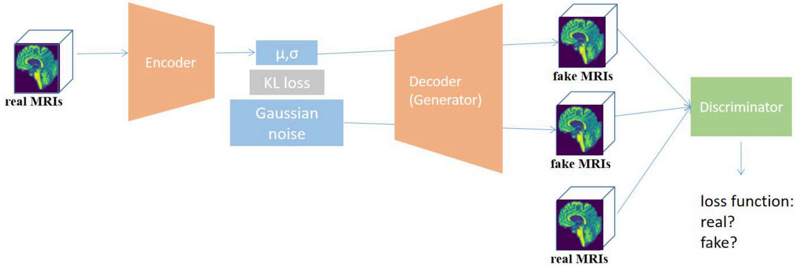

The goal of GAN is to learn a generator G mapping a random vector z_r∈Z (often a mul-tivariate standard normal) to an image X_r∈Xand a discriminator D which differentiates a generated image Xr from the real image. Due to the poor coverage of the Z-space, randomly generated images often suffer from mode collapse. VAE-GAN (Larsen et al., 2016) addresses this issue by learning an encoder E mapping from X to Z to derive the embed-ding size. This results in blurry images, so a recently developed 3D-α-WGAN (Kwon et al.,2019) leverages the CD-based encoder loss from α-GAN (Rosca et al., 2017). Although the mode collapse issue is alleviated in the X-space without blurriness, we identified that E cannot accurately construct the Z-space from 3D brains (e.g., z_e= standard normal prior). We leverage these findings to achieve improved embeddings of 3D brains, which also results in better image quality without mode collapse issue.

3D-α-WGAN

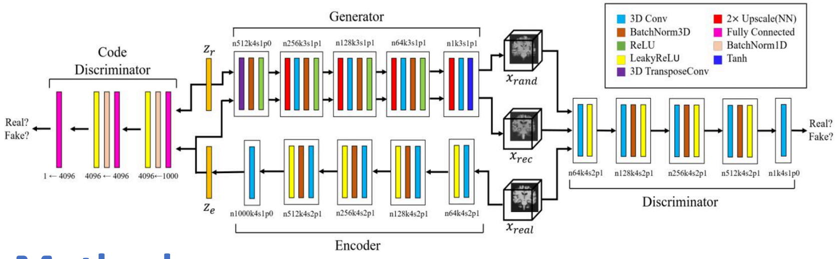

Methods

Our solution builds on 3D-α-GAN. First, instead of using the CD from 3D-α-GAN or KL loss from VAE-GAN, we use the Wasserstein loss explicitly in the Z-space between the random vectors and the embeddings. We refer to this as the Wasserstein Auto-Encoder GAN (WAE-GAN) which provides amore flexible mapping compared to the variational approach while more strictly enforcing the Gaussianness than the CD. Second, we further improve the Z-space by deriving two additional cycle consistent embeddings: z_ee=E(G(z_e)) and z_re=E(G(z_r)). Thus, our final model, Cycle Consistent Embedding GAN (CCE-GAN), solves the following:

$$\underset{D}{\operatorname{argmin}}\ \mathbb{E}_{z_e}[D(X_e)] + \mathbb{E}_{z_r}[D(X_r))] - 2\mathbb{E}_{x_{real}}[D(x_{real})] + \lambda_1 L_{gp}(D)$$

$$\underset{G,E}{\operatorname{argmin}} - \mathbb{E}_{z_e}[D(X_e)] - \mathbb{E}_{z_r}[D(X_r)] + \lambda_2 ||X_r - X_e||_1 + \mathbb{W}_l(z_r, z_e)$$

Cycle Consistent Embedding (CCE) GAN

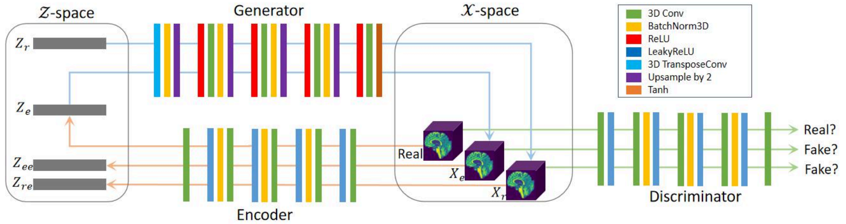

Results

Evaluation: We compute the Maximum Mean Discrepancy (MMD) measure (linear and RBF kernels) between the real images and the generated images for the X-space, and between the random vectors z_rand their corresponding embeddings z_e for the Z-space. We take the average of 100 MMDs. The Structural Similarity (SSIM) measures the distribution diversity where the real data SSIM is 0.839. For each model, we generate 1000 image pairs and compute the average SSIM which aims to be similar to the real data SSIM.

Z-space: We first check the encoder outputs of the real images to evaluatethe embeddings compared to the random standard normal Z-space (1000-D). Fig. 2aand b show the PCA embeddings of 150 random examples. Fig. 2a shows that 3D-α-WGANproduces sparse embeddings, while Fig. 2b shows that VAE-GAN and CCE-GAN produceembeddings highly similar to the random Z-space also shown quantitatively in Table 1.

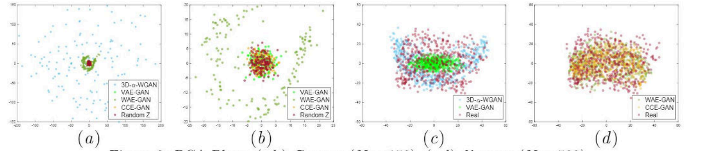

Figure 2: PCA Plots. (a,b) $\mathcal{Z}$-space ($N = 150$), (c,d) $\mathcal{X}$-space ($N = 500$)

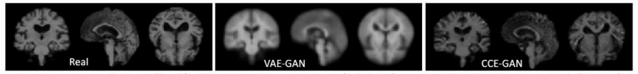

Figure 3: Examples of Real (Left), VAE-GAN generated (Middle), and CCE-GAN generated (Right) brains.

	$\mathcal{X}$-space			$\mathcal{Z}$-space	
Model	**Linear**	**RBF**	**SSIM**	**Linear**	**RBF**
3D-α-GAN	762.4	0.77	**0.841**	619579.4	3.06
VAE-GAN	**354.1**	1.13	0.971	249.4	**0.42**
WAE-GAN	765.4	0.78	0.851	602.1	1.04
CCE-GAN	675.4	**0.73**	0.848	**192.1**	0.59

Table 1: **Linear** and **RBF** MMD, and **SSIM**.

X-space: Table 1 shows the advantage of CCE-GAN, and Fig. 2d shows the PCA embeddings of the images being closer to the real dataset than Fig. 2c. In Fig. 3, wesee VAE-GAN, despite reasonable MMD measures in Z and X, results in blurry images.

Conclusion: We enable accurate mapping of 3D brains in X-space to their embeddingsin the Z-space in a cycle consistent manner using CCE-GAN. We show that if a better embedding is achieved, it also leads to better image generation for 3D MRI generation taskas well. For future work, we will consider other 3D brain datasets.