

Inpainting Cropped Diffusion MRI using Deep Generative Models

Rafi Ayub¹, Qingyu Zhao¹, M. J. Meloy³, Edith V. Sullivan¹, Adolf Pfefferbaum^{1,2}, Ehsan Adeli¹, and Kilian M. Pohl^{1,2}

¹ Stanford University, Stanford, CA, USA

² SRI International, Menlo Park, CA, USA

³ University of California, San Diego, La Jolla, CA, USA

Abstract. Minor artifacts introduced during image acquisition are often negligible to the human eye, such as a confined field of view resulting in MRI missing the top of the head. This cropping artifact, however, can cause suboptimal processing of the MRI resulting in data omission or decreasing the power of subsequent analyses. We propose to avoid data or quality loss by restoring these missing regions of the head via variational autoencoders (VAE), a deep generative model that has been previously applied to high resolution image reconstruction. Based on diffusion weighted images (DWI) acquired by the National Consortium on Alcohol and Neurodevelopment in Adolescence (NCANDA), we evaluate the accuracy of inpainting the top of the head by common autoencoder models (U-Net, VQVAE, and VAE-GAN) and a custom model proposed herein called U-VQVAE. Our results show that U-VQVAE not only achieved the highest accuracy, but also resulted in MRI processing producing lower fractional anisotropy (FA) in the supplementary motor area than FA derived from the original MRIs. Lower FA implies that inpainting reduces noise in processing DWI and thus increase the quality of the generated results. The code is available at <https://github.com/RdoubleA/DWI-inpainting>.

1 Introduction

Diffusion MRI, or diffusion weighted imaging (DWI), is widely used to investigate white matter integrity and structural connectivity between brain regions. Studies based on DWI have revealed disruption in structural networks associated with stroke, brain tumors, neurodegenerative disorders (such as multiple sclerosis), and neuropsychiatric disorders (e.g., schizophrenia) [1]. Despite its wide usage, DWI is plagued by numerous signal artifacts that require extensive pre-processing, such as eddy currents, susceptibility distortion, signal dropout, and motion artifacts [2]. Additionally, part of the brain can be cut off in the image (as shown in Figure 1) due to improper positioning of the subject, limitations in slice acquisitions and prescription, or subject repositioning during a scan session. This cropping artifact can frequently occur in prospective longitudinal studies for investigating neurodevelopment during childhood and adolescence, where imaging acquisition protocols (e.g., field of view) are optimized based on the baseline

visits of subjects and then become suboptimal when the head size increases in later visits. The cropping not only results in missing information, but also influences subsequent preprocessing steps such as inter-subject or cross-modality registration that heavily relies on image boundary information [3]. These misalignments can adversely affect the followup group analysis on regional DWI measures leading to spurious findings.

Correcting for cropped brain regions can be addressed by image inpainting methods [4]. The goal of image inpainting is to predict missing data in corrupted regions based on information from the rest of the image. For example, patch-based methods find candidate replacement patches in the undamaged parts of the image to fill in corrupted regions using matrix-based approaches [5–7] and texture synthesis [8]. Other examples are diffusion-based methods propagating information to corrupted regions from its neighboring areas via interpolation [9, 10]. While these methods have shown promising results, they rely on local image properties, often resulting in reconstructions that ignore global context and thus produce unrealistic looking MRIs [11].

The global context can be learned by deep learning methods trained on large datasets. For example, fully convolutional networks based on the U-Net architecture [12–14] use skip connections to propagate multiscale features to fix the image appearance in the missing regions [15–18]. The U-Net architecture can be further augmented by deep generative models, such as variational autoencoders (VAE) [19]), which first learn a latent distribution explicitly capturing multiscale structures before generating the missing image data. Vanilla VAEs have been implemented for image denoising and inpainting [20], but their low-dimensional latent distribution tend to omit high-frequency features resulting in blurry images [21]. Significant improvement of reconstruction fidelity can be achieved by vector quantized VAEs (VQVAEs) [22, 23]. For example, VQVAE (with skip connections similar to U-Net) have accurately reconstructed T1-weighted structural MRIs [24] but their prediction accuracy on cropped image regions still needs to be tested.

Here, we experiment with several deep generative modeling approaches to repair cropping artifacts in DWI acquired by the longitudinal study National Consortium on Alcohol and Neurodevelopment in Adolescence (NCANDA). Based on DWI that were cropped by us (i.e., the ground truth is known), we compare the inpainting accuracy of U-VQVAE, which combines aspects of U-Net and VQVAE, to U-Net, VQVAE and VAE-GAN [25]. VAE-GAN incorporates an adversarial discriminator network and has previously been applied to image reconstruction. On the real data (i.e., image acquisition caused cropping), we then highlight the improvement of processing based on the inpainted DWIs by computing fractional anisotropy (FA) in regions affected by the cropping. Compared to the regional FA from the original DWI, FA derived based on the inpainted DWI were lower, indicating a reduction in noise in the image processing pipeline. These results demonstrate the utility of autoencoder models for repairing cropping artifacts in multishell diffusion MRI and improving signal-to-noise ratio in downstream processing steps. More important, they also provide us with a tool

for restoring parts of the brain in MRIs that are cut off due to improper subject positioning or suboptimal field-of-view.

2 Methods

Since the U-Net architecture has been extensively discussed in prior studies, we focus on describing the U-VQVAE architecture. From that, VQVAE and VAE-GAN are derived. After that we introduce the dataset used for training the models and the metrics used for evaluation.

2.1 Model Architecture

U-VQVAE reconstructed cropped DWI by leveraging the generative modeling capabilities of the VAE architecture. VAE [19] consists of an encoder ϕ and a decoder network ψ . The encoder models a posterior distribution $p(\mathbf{z}|\mathbf{x})$ of the latent random variable \mathbf{z} given the input DWI image \mathbf{x} , and the decoder models $p(\mathbf{x}|\mathbf{z})$, the likelihood of input image given \mathbf{z} . Typically, \mathbf{z} is assumed to be normally distributed $\mathcal{N}(0, 1)$, but VQVAEs [22, 23] model \mathbf{z} with a collection of embedding vectors $\mathbf{E} \in R^{K \times D}$, where K is the predefined number of embedding vectors and D is the predefined dimensionality of each embedding vector $\mathbf{e}_k \in R^D, k \in \{1, 2, \dots, K\}$. In our model, we set $K = 512$ and $D = 32$. We assume that the encoder first reduces the input \mathbf{x} to a 3D volume of J voxels before entering the latent space, with the j^{th} voxel characterized by a D -dimensional vector \mathbf{z}_e^j (D channels). The encoder ϕ then further incorporates a quantization that maps each \mathbf{z}_e^j to the nearest vector in the embedding space $\mathbf{E} = \{\mathbf{e}_0, \dots, \mathbf{e}_{K-1}\}$. Then the final embedded latent representation \mathbf{z}^j follows a categorical posterior distribution:

$$q(\mathbf{z}^j = \mathbf{e}_\tau | \mathbf{x}) = \begin{cases} 1 & \text{for } \tau = \operatorname{argmin}_k \|\mathbf{z}_e^j - \mathbf{e}_k\|_2 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

In other words, quantization ensures that the decoder uses discrete vectors from the embedding for a high-resolution reconstruction instead of a fuzzy encoding created by the encoder alone. The embedding space is not predetermined, but instead is learned through back propagation with exponential moving average as described in [22].

The loss function for the U-VQVAE model consisted of an image reconstruction loss, a codebook loss, and a commitment loss. The image reconstruction loss is defined by the Mean Squared Error between the input of the encoder \mathbf{x} , i.e, the cropped MRI, and the ground-truth \mathbf{x}' , i.e, the MRI without cropping. In other words the loss is $\|\mathbf{x}' - \psi(\phi(\mathbf{x}))\|_2$, which is dependent on both the encoder and decoder networks. The codebook loss uses the l_2 error to encourage the embedding vectors in \mathbf{E} to move closer towards the learned features \mathbf{z}_e^j ; i.e, it ensures that the discrete embedding accurately reflects the compressed representation of the image learned by the encoder. Defining $\operatorname{sg}(\cdot)$ as the stop-gradient operation

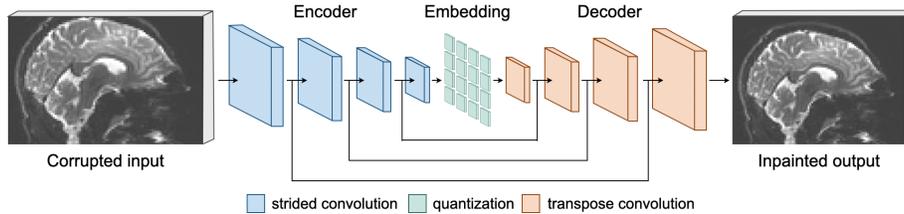


Fig. 1: Architecture of the U-VQVAE model

and \mathbf{e}_τ^j as the final embedded vector that matches \mathbf{z}_e^j (Eq. 1), this loss term is defined by $\sum_{j=1}^J \|\mathbf{sg}(\mathbf{z}_e^j) - \mathbf{e}_\tau^j\|_2^2$. The commitment loss also relies on the l_2 error to encourage \mathbf{z}_e^j to converge to the embedding vectors. This loss term is defined by $\sum_{j=1}^J \|\mathbf{z}_e^j - \mathbf{sg}(\mathbf{e}_\tau^j)\|_2^2$, which is only dependent on the encoder. Let $\beta = 6$ be the weight associated with the commitment loss and $\alpha = 1$ the weight of the codebook loss, the objective function associated with an input image \mathbf{x} is then defined as

$$\mathcal{L}(\mathbf{x}, \mathbf{E}) := \|\mathbf{x}' - \psi(\phi(\mathbf{x}))\|_2 + \sum_{j=1}^J (\alpha \|\mathbf{z}_e^j - \mathbf{sg}(\mathbf{e}_\tau^j)\|_2^2 + \beta \|\mathbf{z}_e^j - \mathbf{sg}(\mathbf{e}_\tau^j)\|_2^2) \quad (2)$$

The encoder network of the model (see also Figure 1) consists of four convolutional layers, where each layer downsamples the DWI by a factor of two. The decoder network of the model is defined by four transpose convolutional layers, where each layer upsamples the compressed representation by a factor of two. ReLU activations are used in every layer. All convolutional layers have kernel size $4 \times 4 \times 4$, a stride of 2, and padding of 1.

With regard to the other three models, the U-Net model is implemented as in [13] except with 4 filters in the first convolutional layer as opposed to 64. This modification can effectively reduce the model size so that the batch size can be increased. U-Net uses only the reconstruction loss without the codebook and commitment loss. The VQVAE model removes the skip connections on every level in the U-VQVAE. The VAE-GAN model uses the same U-VQVAE architecture with additional batch normalization layers as a generator and adds a fully convolutional patch-wise discriminator classifier [14] in an adversarial setup. Additionally, it produces less blurry reconstructions by replacing the voxel-wise l_2 reconstruction loss with the l_2 error between real and generated image representations in the third convolutional layer of the discriminator network [25].

2.2 Dataset and Evaluation Metrics

We first trained and tested the model on an artificially cropped dataset, which was based on the b0 and b1000 DWIs of 824 subjects from NCANDA without cropping artifacts (Public Release: NCANDA_PUBLIC_BASE_DIFFUSION_V01) and of 100 subjects from the Human Connectome Project (HCP) [26,27]. Images

Model	Whole				Cropped			
	SSIM	PSNR	MSE	Gradient	SSIM	PSNR	MSE	Gradient
U-VQVAE	0.9973	49.8733	1.089e-5	0.0209	0.9679	30.7809	8.535e-5	0.0041
VQVAE	0.8467	32.0512	5.437e-4	0.0064	0.9279	28.4808	1.449e-4	0.0034
U-Net [13]	0.9959	49.4796	1.152e-5	0.0210	0.9638	30.6117	8.723e-5	0.0039
VAE-GAN	0.8074	30.6311	2.919e-3	0.1206	0.8036	21.7169	2.348e-3	0.2033

Table 1: Model performance evaluated by various image quality metrics on both the entire image and only the cropped regions. Amount of cropping was fixed to 12.5%, similar to the training data. PSNR is measured in dB.

from HCP were downsampled to $96 \times 96 \times 64$ to match the resolution of NCANDA images. All images were normalized to the range 0-1 to minimize influence from variation in voxel intensities in the training the models. The training data was further augmented by randomly rotating images ten degrees in either direction for every axis and translating 10 voxels in the two directions of the axial plane. This resulted in a dataset of 420,487 MRI, 236,291 MRIs from NCANDA and 184,196 MRIs from HCP. The MRIs consisted of 93,467 b0 volumes and 327,020 b1000 volumes. 80% of the dataset was used for training and 20% for testing.

The optimal model was determined by its capability of inpainting DWI cropped by us. Specifically, training MRIs were cropped eight slices (or 12.5%) from the top before being analyzed by the model. This amount of cropping reflected the average cropping across the real corrupted MRIs of the NCANDA dataset. Models were trained on a single NVIDIA Tesla V100-SXM2 32 GB GPU. The Adam optimizer was used with a learning rate of 0.001. The batch size was set to 128 DWI images for all models. To measure the inpainting accuracy, we applied the models to the artificially cropped MRIs of the testing data and compared the reconstruction to the ground-truth via Structural Similarity [28], Peak Signal-to-Noise Ratio (PSNR), and Mean-Squared Error (MSE). We also computed the average image spatial gradient as an indicator of how well models recreated high-frequency details. Separately, these metrics were calculated confined to the eight cropped slices.

Using the optimal model derived by these metrics, we evaluated the impact of inpainting on downstream processing steps. The experiment was based on the DWI sequences of 13 adolescents at their followup visits (from Year 2 to Year 6) of the NCANDA dataset. These DWIs were labelled as ‘unusable’ by the NCANDA quality assurance protocol [29] due to the field of view not being able to capture the entire head anymore (top of the brain is cut off, see also Fig. 1) as it increased in size compared to the baseline. To show the impact of inpainting, we registered the FA maps of the DWI to the SRI24 atlas space [30] based on either the real b0 volumes with cropping artifacts or the inpainted volumes. We then analyzed the differences in FA in between those two FA maps to evaluate the influence from the inpainting.

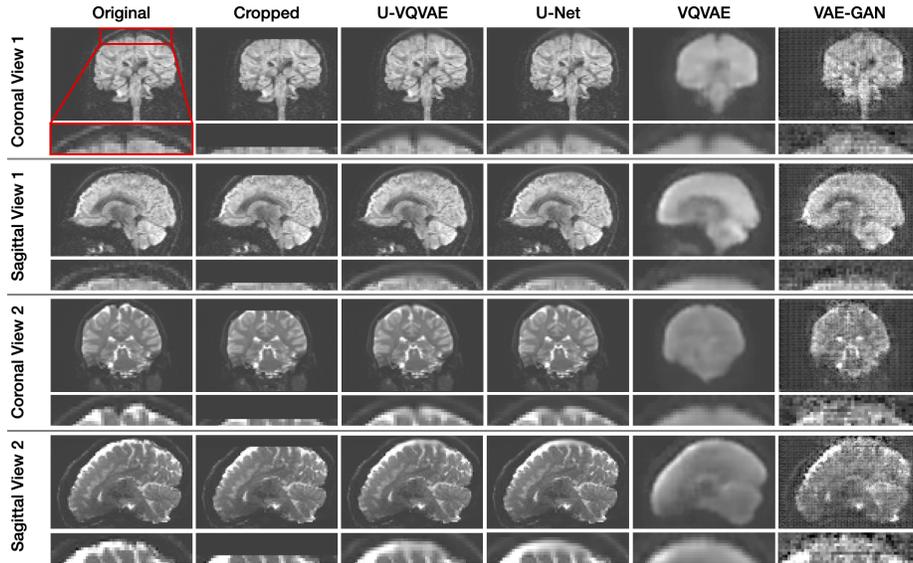


Fig. 2: Example b0 and non-b0 original and inpainted images for each model. Each row is a different view of the image with the bottom half zoomed in on the cropped region.

3 Results and Discussion

3.1 Evaluation on artificially cropped B0 MRIs of DWI

The accuracy scores for each model on the artificially cropped dataset are listed in Table 1. While VAE-GAN achieved the highest mean image gradient among the models, it also recorded the lowest PSNR. This is likely due to the relatively high level of noise levels shown in the inpainted images from this model in Figure 2 (compared to those from U-VQVAE). The increased noise may have been introduced by adversarial training of the generator with the discriminator network, which may have conflicted with reconstruction of intact slices of the image through the generator’s skip connections. Across all other metrics, U-VQVAE performed the best. As expected, image quality metrics for U-VQVAE were much better than those for VQVAE, indicating the importance of including skip connections. This is also supported by U-Net’s high SSIM and PSNR, which nearly rivals U-VQVAE. To determine whether U-VQVAE had statistically higher reconstruction fidelity than U-Net, we performed a two-sample t-tests on U-Net’s and U-VQVAE’s mean SSIM and PSNR for both whole image and cropped parts. U-VQVAE’s increase in accuracy compared to U-Net was statistically significant in all four cases (i.e, $p < 0.001$; p-values were one-sided and FDR corrected). These results demonstrate that U-VQVAE was the most accurate model for reconstructing cropped regions informed by the image’s global structure.

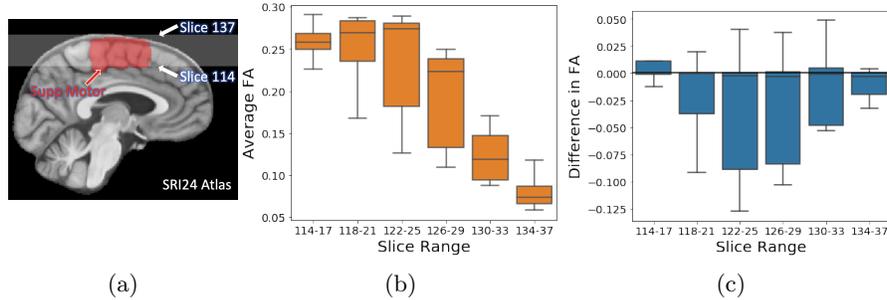


Fig. 3: (a) The Supplementary Motor Area overlaid with the SRI24 atlas; (b) Average FA at different axial slices of the 13 NCANDA DWI sequences with cropping artifacts; (c) Difference in average FA derived by registrations based on either cropped or inpainted b0 volumes

3.2 Impact on Downstream Preprocessing

This section highlights the potential usage of our proposed inpainting method by U-VQVAE in analyzing NCANDA data, which in part studies the white matter microstructural development during adolescence. The DWI sequence of each of the 13 subjects used in this experiment was preprocessed by the publicly available NCANDA image processing pipeline [29], which included bad single shots removal, echo-planar structural distortion correction, eddy-current correction, rigid alignment for motion correction, and skull stripping. The Fractional Anisotropy (FA) map was estimated by CAMINO [31].

Performing population-level analysis on the FA maps requires aligning them to a single template. The alignment was performed by non-rigidly registering the b0 volume to the subject T2w image, which was then non-rigidly registered to the SRI24 atlas [30] (Fig. 3a). This alignment was potentially corrupted at the top of the skull, specifically in the Supplementary Motor Area (SMA, Fig. 3a), due to the cropping artifact. We show this by repeating the registration using the inpainted b0 volumes. Fig. 3c shows the difference in the average FA of SMA at different axial slices (slice 114 to 137) derived by the two registration approaches. We can see that while the two registrations resulted in similar FA in lower axial slices (slice < 118) and outside the brain (slice > 134), the registration based on inpainted b0 volumes generated lower FA at slice 118 to 133. This difference was more pronounced at slice 122 to 129, regions most severely impacted by the cropping, despite the average FA value decreased in magnitude at these slices (Fig. 3b). This indicated that our inpainting could potentially improve the power of group-level analysis on the FA measures as higher FA estimates are generally associated with greater noise or larger artifacts within the data [32].

4 Conclusion

In this study, we presented a vector quantized variational autoencoder architecture that is capable of repairing cropping artifacts in multi-shell diffusion weighted images for the first time. The images inpainted by our model exhibited higher fidelity as measured by various image quality metrics than other landmark models widely used in image reconstruction. Most importantly, images inpainted by our model yielded FA maps with lower FA values in areas previously impacted by the cropping artifact, indicating a reduction in noise in the image processing pipeline and demonstrating that our inpainting can improve the power of group-level analyses. Future directions of our modeling approach could include training with varied levels of cropping to improve model robustness and generalizing the model to other MR imaging modalities.

Acknowledgements. This work was supported by NIH Grants AA021697, AA005965, and AA010723. This work was also supported by the National Science Foundation Graduate Research Fellowship and the 2020 HAI-AWS Cloud Credits Award.

References

1. Soares, J.M., Marques, P., Alves, V., Sousa, N.: A hitchhiker’s guide to diffusion tensor imaging. *Frontiers in Neuroscience* **7**(7 MAR), 1–14 (2013). <https://doi.org/10.3389/fnins.2013.00031>
2. Le Bihan, D., Poupon, C., Amadon, A., Lethimonnier, F.: Artifacts and pitfalls in diffusion MRI. *Journal of Magnetic Resonance Imaging* **24**(3), 478–488 (2006). <https://doi.org/10.1002/jmri.20683>
3. Greve, D.N., Fischl, B.: Accurate and robust brain image alignment using boundary-based registration. *NeuroImage* **48**(1), 63–72 (oct 2009). <https://doi.org/10.1016/j.neuroimage.2009.06.060>, <https://linkinghub.elsevier.com/retrieve/pii/S1053811909006752>
4. Elharrouss, O., Almaadeed, N., Al-Maadeed, S., Akbari, Y.: Image Inpainting: A Review. *Neural Processing Letters* **51**(2), 2007–2028 (apr 2020). <https://doi.org/10.1007/s11063-019-10163-0>, <http://link.springer.com/10.1007/s11063-019-10163-0>
5. Lu, H., Liu, Q., Zhang, M., Wang, Y., Deng, X.: Gradient-based low rank method and its application in image inpainting. *Multimedia Tools and Applications* **77**(5), 5969–5993 (2018). <https://doi.org/10.1007/s11042-017-4509-0>
6. Jin, K.H., Ye, J.C.: Annihilating Filter-Based Low-Rank Hankel Matrix Approach for Image Inpainting. *IEEE Transactions on Image Processing* **24**(11), 3498–3511 (2015). <https://doi.org/10.1109/TIP.2015.2446943>
7. Guo, Q., Gao, S., Zhang, X., Yin, Y., Zhang, C.: Patch-Based Image Inpainting via Two-Stage Low Rank Approximation. *IEEE Transactions on Visualization and Computer Graphics* **24**(6), 2023–2036 (2018). <https://doi.org/10.1109/TVCG.2017.2702738>
8. Kozhikin, N., Savchenko, V., Senin, M., Hagiwara, I.: An approach to surface retouching and mesh smoothing. *Visual Computer* **19**(7-8), 549–564 (2003). <https://doi.org/10.1007/s00371-003-0218-y>

9. Chan, T.F., Shen, J.: Nontexture inpainting by curvature-driven diffusions. *Journal of Visual Communication and Image Representation* **12**(4), 436–449 (2001). <https://doi.org/10.1006/jvci.2001.0487>
10. Alsalamah, M., Amin, S.: Medical Image Inpainting with RBF Interpolation Technique. *International Journal of Advanced Computer Science and Applications* **7**(8), 91–99 (2016). <https://doi.org/10.14569/ijacsa.2016.070814>
11. Yan, Z., Li, X., Li, M., Zuo, W., Shan, S.: Shift-net: Image inpainting via deep feature rearrangement. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **11218 LNCS**, 3–19 (2018). https://doi.org/10.1007/978-3-030-01264-9_1
12. Shelhamer, E., Long, J., Darrell, T.: Fully Convolutional Networks for Semantic Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **39**(4), 640–651 (2017). <https://doi.org/10.1109/TPAMI.2016.2572683>
13. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **9351**, 234–241 (2015). https://doi.org/10.1007/978-3-319-24574-4_28
14. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017 2017-Janua*, 5967–5976 (2017). <https://doi.org/10.1109/CVPR.2017.632>
15. Armanious, K., Mecky, Y., Gatidis, S., Yang, B.: Adversarial Inpainting of Medical Image Modalities. In: *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 3267–3271. IEEE (may 2019). <https://doi.org/10.1109/ICASSP.2019.8682677>, <https://ieeexplore.ieee.org/document/8682677/>
16. Armanious, K., Kumar, V., Abdulatif, S., Hepp, T., Gatidis, S., Yang, B.: ipA-MedGAN: Inpainting of Arbitrary Regions in Medical Imaging (2019), <http://arxiv.org/abs/1910.09230>
17. Armanious, K., Gatidis, S., Nikolaou, K., Yang, B., Kustner, T.: Retrospective correction of rigid and non-rigid mr motion artifacts using gans. *Proceedings - International Symposium on Biomedical Imaging 2019-April*, 1550–1554 (2019). <https://doi.org/10.1109/ISBI.2019.8759509>
18. Sabokrou, M., Pourreza, M., Fayyaz, M., Entezari, R., Fathy, M., Gall, J., Adeli, E.: AVID: Adversarial Visual Irregularity Detection. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **11366 LNCS**, 488–505 (2019). https://doi.org/10.1007/978-3-030-20876-9_31
19. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *2nd International Conference on Learning Representations, ICLR 2014 - Conference Track Proceedings (MI)*, 1–14 (2014)
20. Xie, J., Xu, L., Chen, E.: Image denoising and inpainting with deep neural networks. *Advances in Neural Information Processing Systems* **1**, 341–349 (2012)
21. Dosovitskiy, A., Brox, T.: Generating images with perceptual similarity metrics based on deep networks. *Advances in Neural Information Processing Systems* pp. 658–666 (2016)
22. Van Den Oord, A., Vinyals, O., Kavukcuoglu, K.: Neural discrete representation learning. *Advances in Neural Information Processing Systems 2017-Decem(Nips)*, 6307–6316 (2017)
23. Razavi, A., van den Oord, A., Vinyals, O.: Generating Diverse High-Fidelity Images with VQ-VAE-2 (2019), <http://arxiv.org/abs/1906.00446>

24. Tudosiu, P.D., Varsavsky, T., Shaw, R., Graham, M., Nachev, P., Ourselin, S., Sudre, C.H., Cardoso, M.J.: Neuromorphologically-preserving Volumetric data encoding using VQ-VAE pp. 1–13 (2020), <http://arxiv.org/abs/2002.05692>
25. Larsen, A.B.L., Sønderby, S.K., Larochelle, H., Winther, O.: Autoencoding beyond pixels using a learned similarity metric. 33rd International Conference on Machine Learning, ICML 2016 **4**, 2341–2349 (2016)
26. Van Essen, D.C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T.E., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S.W., Della Penna, S., Feinberg, D., Glasser, M.F., Harel, N., Heath, A.C., Larson-Prior, L., Marcus, D., Michalareas, G., Moeller, S., Oostenveld, R., Petersen, S.E., Prior, F., Schlaggar, B.L., Smith, S.M., Snyder, A.Z., Xu, J., Yacoub, E.: The Human Connectome Project: A data acquisition perspective. *NeuroImage* **62**(4), 2222–2231 (2012). <https://doi.org/10.1016/j.neuroimage.2012.02.018>
27. Hodge, M.R., Horton, W., Brown, T., Herrick, R., Olsen, T., Hileman, M.E., McKay, M., Archie, K.A., Cler, E., Harms, M.P., Burgess, G.C., Glasser, M.F., Elam, J.S., Curtiss, S.W., Barch, D.M., Oostenveld, R., Larson-Prior, L.J., Ugurbil, K., Van Essen, D.C., Marcus, D.S.: ConnectomeDB—Sharing human brain connectivity data. *NeuroImage* **124**(3), 1102–1107 (jan 2016). <https://doi.org/10.1016/j.neuroimage.2015.04.046>, <https://linkinghub.elsevier.com/retrieve/pii/S1053811915003468>
28. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**(4), 600–612 (2004). <https://doi.org/10.1109/TIP.2003.819861>
29. Pohl, K.M., Sullivan, E.V., Rohlfing, T., Chu, W., Kwon, D., Nichols, B.N., Zhang, Y., Brown, S.A., Tapert, S.F., Cummins, K., Thompson, W.K., Brumback, T., Colrain, I.M., Baker, F.C., Prouty, D., De Bellis, M.D., Voyvodic, J.T., Clark, D.B., Schirda, C., Nagel, B.J., Pfefferbaum, A.: Harmonizing DTI measurements across scanners to examine the development of white matter microstructure in 803 adolescents of the NCANDA study. *NeuroImage* **130**, 194–213 (2016). <https://doi.org/10.1016/j.neuroimage.2016.01.061>
30. Rohlfing, T., Zahr, N.M., Sullivan, E.V., Pfefferbaum, A.: The SRI24 multichannel atlas of normal adult human brain structure. *Human Brain Mapping* **31**(5), 798–819 (2010). <https://doi.org/10.1002/hbm.20906>
31. Cook, P.a., Bai, Y., Seunarine, K.K., Hall, M.G., Parker, G.J., Alexander, D.C.: Camino: Open-Source Diffusion-MRI Reconstruction and Processing. 14th Scientific Meeting of the International Society for Magnetic Resonance in Medicine **14**, 2759 (2006)
32. Farrell, J.A.D., Landman, B.A., Jones, C.K., Smith, A., Prince, J.L., Zijl, P.C.M.V., Mori, S.: Effects of SNR on the Accuracy and Reproducibility of DTI-derived Fractional Anisotropy, Mean Diffusivity, and Principal Eigenvector Measurements at 1.5T. *Journal of Magnetic Resonance* **26**(3), 756–767 (2010). <https://doi.org/10.1002/jmri.21053.Effects>