

3D Brain MRI GAN-based synthesis conditioned on Partial Volume Maps

Filip Rusak^{1,2(✉)}, Rodrigo Santa Cruz^{1,2}, Pierrick Bourgeat², Clinton Fookes¹, Jurgen Fripp², Andrew Bradley¹, and Olivier Salvado^{1,2}

¹ Queensland University of Technology, Brisbane 4000, QLD, Australia

² CSIRO, Herston QLD 4029, Australia

filip.rusak@data61.csiro.au

Abstract. In this paper, we propose a framework for synthesising 3D brain T1-weighted (T1-w) MRI images from Partial Volume (PV) maps for the purpose of generating synthetic MRI volumes with more accurate tissue borders. Synthetic MRIs are required to enlarge and enrich very limited data sets available for training of brain segmentation and related models. In comparison to current state-of-the-art methods, our framework exploits PV-map properties in order to guide a Generative Adversarial Network (GAN) towards the generation of more accurate and realistic synthetic MRI volumes. We demonstrate that conditioning a GAN on PV-maps instead of Binary-maps results in 58.96% more accurate tissue borders in synthetic MRIs. Furthermore, our results indicate an improvement in the representation of the Deep Gray Matter region in synthetic MRI volumes. Finally, we show that fine changes introduced into PV-maps are reflected in the synthetic images, while preserving accurate tissue borders, thus enabling better control during the data synthesis of novel synthetic MRI volumes.

Keywords: Generative Adversarial Network · Partial Volume Maps · Synthetic MRIs · 3D Image Synthesis.

1 Introduction

Deep Neural Networks, particularly Convolutional Neural Networks (CNNs), have demonstrated tremendous capability to perform accurate segmentation tasks when trained on large datasets [19,20]. In medical imaging, these methods are limited by the scarcity of available data. Labelling medical data is time consuming and requires a high level of expertise which is expensive. Many different CNN-based methods attempted to overcome this hurdle by mitigating the amount of data needed for their training, such as using unsupervised [3,15], weakly-supervised [8,27], semi-supervised [2,17] and self-supervised [14,21] methods. The drawback of these methods is that they are typically less accurate than

Preprint, final version published in International Workshop on Simulation and Synthesis in Medical Imaging (pp. 11-20). Springer, Cham, 2020. The final authenticated publication is available online at https://doi.org/10.1007/978-3-030-59520-3_2.

supervised methods [13]. Furthermore, since the ground truth label is missing, it is more difficult to evaluate the performance of these methods [12].

In contrast to aforementioned methods, data augmentation methods [5,22,23] aim to increase the number of available labelled samples needed for training of supervised methods. Data augmentation methods fall into two major trends: geometric transformation-based and GAN-based. Most geometric transformation-based augmentation methods provide limited improvement in terms of samples variety as their output highly relies on the input data.

A GAN is a data synthesis approach capable of injecting more variety into synthesised data and generating outputs less dependant of the input data, while aiming to follow the training data distribution [6]. MRI synthesis using GANs can be classified into two prominent approaches: unconditional [6,7,11] and conditional [16,22]. The main drawback of unconditional MRI synthesis approaches, in the context of supervised segmentation, is the missing segmentation labels of the newly synthesised MRIs. Another drawback of such approaches is the lack of synthesis control [16]. On the other hand, MRI synthesis approaches based on conditioning a GAN with segmentation labels, as presented in [22], keeps the brain anatomical structures intact, while segmentation labels give the ability to control the synthetic results. Nevertheless, the segmentation labels only provide an estimate of brain tissue types. Their accuracy is limited by the image resolution and consequently the segmentation accuracy may suffer from partial volume (PV) effects at the border between two tissues where a single voxel may contain multiple classes. More accurate segmentation can be represented with PV-maps as they define accurate border between two tissue classes [4], which makes them a suitable choice for conditioning GANs in the context of MRI synthesis. Conditioning GANs on PV-maps opens a pathway to generate MRIs of different appearances while retaining the same anatomical structure with fine boundary details. Having control over MRI synthesis by defining tissues with PV-maps as well as the ability to change them may be used as a powerful data synthesis approach.

In this paper, we propose a framework for synthesising 3D brain T1-weighted MRI images from PV-maps. Our proposed framework is inspired by well-known Image-to-Image conditional GAN approach described in [9]. We use PV-maps of Gray Matter (GM), White Matter (WM) and Cerebrospinal Fluid (CSF) as inputs to assist the training of the model and the generation of realistic 3D brain MRIs. We report the first attempt to synthesise realistic 3D brain MRI images from PV-maps using GANs. Furthermore, we demonstrate that changes in PV-maps reflect changes in newly generated synthetic images and show how the framework can increase the number of synthetic training images. The contributions of this paper are the following:

- 1) *We proposed a GAN-based framework that exploits PV-map properties to obtain synthetic MRI volumes with accurate borders between tissue classes as well as more accurate and realistic Deep Gray Matter (DGM) regions.*
- 2) *In the context of 3D T1-w brain MRI generation using GANs, we demonstrated that conditioning GANs on PV-maps produces better results than binary-*

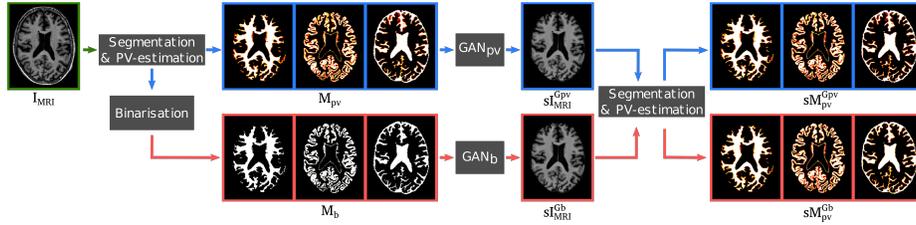


Fig. 1. Schematic representation of the experimental method.

maps. The difference is most evident in the regions of tissue borders, which is an important feature for applications such as cortical thickness estimation and segmentation.

2 Methods

Hypothesis Formulation. When it comes to T1-w brain MRI synthesis, a desirable synthetic MRI image (sI_{MRI}) is expected to respect relations between brain anatomical structures of the original MRI images (I_{MRI}). A possible method to generate such images is to condition a GAN on a particular class label to obtain results that meet the imposed condition [16]. The same mechanism may be applied to the problem of generating sI_{MRI} that keeps the brain anatomy intact. One of the simple ways is to condition a GAN for the purpose of sI_{MRI} generation with intact anatomy is to use Binary Maps (M_b) of different tissues. A M_b , in the context of 3D images, is a volume $M_b \in \{0, 1\}^{w \times h \times d}$, where the value of each voxel denotes affiliation to a single class (1 indicates class affiliation). In the case of brain synthesis, a GAN can be conditioned on three classes: WM, GM and CSF; where each class is represented as a M_b . The limitation of such a class labelling method is the indivisible nature of voxel affiliation. In certain regions of an MRI, especially in the region around a tissue border, the voxel may not be of an adequate size. This limitation can be overcome by using PV-maps (M_{pv}) which, in the context of 3D images, is defined as a volume $M_{pv} \in [0, 1]^{w \times h \times d}$, where the value of each voxel represents the proportion of affiliation to a single class (1 indicates 100% class affiliation). The main advantage of M_{pv} is the ability to represent partial affiliation to a certain class, which allows tissue labelling with higher precision when compared to single-class voxels.

We hypothesise that conditioning a GAN with M_{pv} instead of M_b results with better sI_{MRI} , especially at tissue interfaces. The hypothesis was evaluated by the experimental method presented in Fig. 1. The Fig. 1 shows the generation of M_{pv} from I_{MRI} by performing brain segmentation, implemented with the Expectation-maximisation (EM) algorithm [25], followed by PV-estimation implemented as in [1]. Three M_{pv} s are derived from I_{MRI} , one for each tissue-type (WM, GM and CSF). We binarise M_{pv} s by assigning each voxel to the M_{pv} with the highest partial affiliation for a particular voxel and obtain the corresponding M_b for each class. Two models were trained, GAN_{pv} on M_{pv} and GAN_b on

M_b and used to generate synthetic images, sI_{MRI}^{Gpv} and sI_{MRI}^{Gb} respectively. Once the sI_{MRI} were synthesised, the reverse process was performed, where sI_{MRI}^{Gpv} and sI_{MRI}^{Gb} were segmented followed by PV-estimation in order to obtain the synthetic M_{pv} (sM_{pv}). sM_{pv} derived from sI_{MRI}^{Gb} are denoted as sM_{pv}^{Gb} , while sM_{pv} derived from sI_{MRI}^{Gpv} are denoted as sM_{pv}^{Gpv} . We generated sM_{pv} in order to evaluate to what extent are the imposed conditions preserved in sI_{MRI} .

Model Architecture. The architecture of our model was inspired by Pix2Pix [9] and adapted to facilitate the needs of 3D MRI images. Pix2Pix is a conditional GAN capable of translating labels into images that follow a certain distribution, which makes it suitable for many image-to-image translation problems. The network is composed of a U-net-based generator [18] and a PatchGAN-based discriminator that compares image patches instead of whole images [9]. The modified architecture and its hyper-parameters are presented in Fig. 2.

We denote data of a certain distribution d_x with x , generator with G , its output $G(c_{1-3}, z)$ and discriminator with D . Moreover, we denote three condition variables with $c_{1-3}(M_b$ or M_{pv} for three tissue-types) and a noise variable with z . The objective function is defined as follows,

$$\min_G \max_D \mathbb{E}_{c_{1-3}, x} \left[\log(D(c_{1-3}, x)) \right] + \mathbb{E}_{c_{1-3}, z} \left[\log(1 - D(c_{1-3}, G(c_{1-3}, z))) \right] + \mathbb{E}_{c_{1-3}, x, z} \left[\|x - G(c_{1-3}, z)\|_1 \right], \quad (1)$$

where G has a goal to minimise the probability of D performing a correct binary classification task, while D aims to maximise the same. Referring to [9], we also added the L1 distance clause to the objective function as L1 tends to mitigate blurriness in the resulting images, which is needed for generation of images with accurate tissue borders. We also used the noise z in the form of dropout (activated at training and inference) across a number of layers instead of providing it as an input.

Data. For the evaluation of our training method we used a subset of 3T scans (181x218x181 voxels) from the ADNI [10,26] dataset. The subset contained 700 baseline subjects where only 3D T1-w volumes were used. Subjects were split into train and test sets. The train set included 500 subjects, while the 200 remaining subjects were used for the test set. All volumes were pre-processed by applying: (i) bias field correction in the brain region [24], (ii) rigid registration to the MNI-

Data used in the preparation of this article were obtained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). The ADNI was launched in 2003 as a public-private partnership, led by Principal Investigator Michael W. Weiner, MD. The primary goal of ADNI has been to test whether serial magnetic resonance imaging (MRI), positron emission tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of mild cognitive impairment (MCI) and early Alzheimer’s disease (AD). For up-to-date information, see www.adni-info.org.

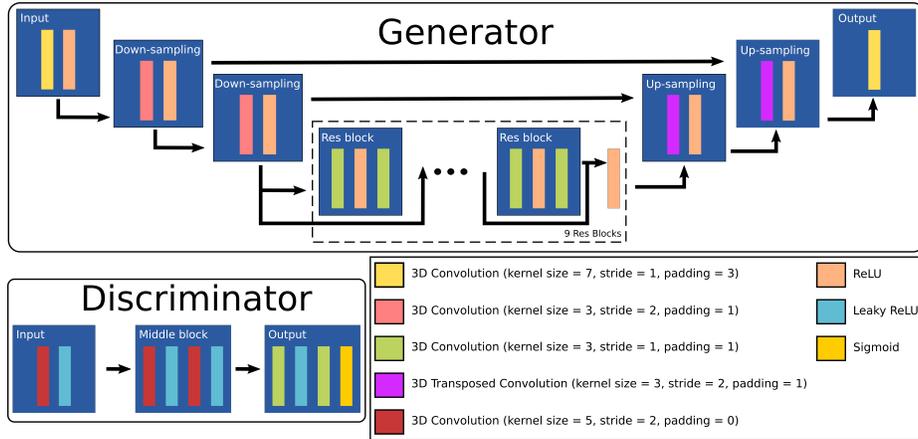


Fig. 2. Model architecture with supplementary hyper-parameter details.

space and (iii) zero-mean normalisation with the mean value computed from the voxels in brain region of interest (ROI) only.

Training. We trained our models for 200 epochs. For the training of both models we used Adam optimiser, batch size of 1 and initial learning rate of 0.0002. After 100 epochs, we reduced the learning rate by 2×10^{-6} every epoch.

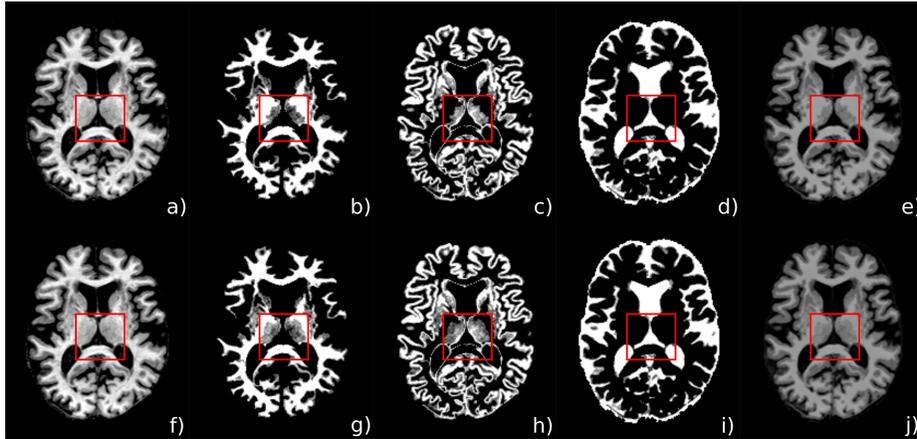


Fig. 3. Qualitative results of our framework trained on M_{pv} . Presented results show that changes introduced in the M_{pv} are reflected in the $sMRI$. The ground truth (a,f), two sets of M_{pv} from the same subject as well as corresponding $sMRI$ (e,j) are presented respectively, where the region of DGM in the first case (b,c,d) is weakly defined, in comparison to the second case (g,h,i).

Table 1. Metrics computed between I_{MRI} and sI_{MRI} created by GAN_b and GAN_{pv} .

	PSNR	MAE	MSE	SSIM
GAN_b	32.777 ± 1.041	0.166 ± 0.024	0.054 ± 0.014	0.955 ± 0.01
GAN_{pv}	33.449 ± 1.103	0.144 ± 0.023	0.047 ± 0.013	0.96 ± 0.01

3 Experiments

Our experiments were constructed to assess the benefit of using M_{pv} over M_b for the purpose of synthesising T1-w brain MRI volumes with accurate tissue-borders. Moreover, as a proof of concept for MRI synthesis, we assessed the reflection of fine changes, introduced on the M_{pv} , in sI_{MRI} and sM_{pv} . In the following experiments we evaluated the quality of sI_{MRI} and sM_{pv} on the level of the brain volume, three tissue ROIs, tissue borders and the region of DGM.

Image Synthesis Quality. We evaluated our models by generating sI_{MRI} from both, M_b and M_{pv} , and comparing them with the corresponding I_{MRI} . Images were compared by employing the following metrics: Peak Signal-to-Noise Ratio (PSNR), Mean Absolute Error (MAE), Mean Squared Error (MSE) and Structural Similarity (SSIM) (see quantitative results in Table 1). PSNR, MSE and MAE were computed in the brain ROI. The dynamic range measured in the brain ROI of I_{MRI} spans between $[-0.56, 9.88]$, and was used to compute PSNR. SSIM was calculated on the whole volume, with background values set to zero as our generator generates brain sI_{MRI} without a background. Table 1 shows that GAN_{pv} produced sI_{MRI} more similar to I_{MRI} than GAN_b .

Evaluation at tissue level. We took a closer look and evaluated the quality of sI_{MRI} as well as the corresponding sM_{pv} in the ROI for every tissue-class (WM, GM and CSF). The GAN , segmentation and PV estimation may introduce errors in either sI_{MRI} or sM_{pv} . Therefore, we computed MAE and MSE between I_{MRI} and sI_{MRI} in order to evaluate the error introduced by GAN . We also used the Dice similarity metric (DSM) to evaluate the overlap with the ground truth and MAE as well as MSE to evaluate the error in sM_{pv} introduced by GAN , segmentation and PV estimation. Quantitative results of the error metrics for each tissue type, calculated on sI_{MRI} , are presented in Table 2. Quantitative measurements of shape and intensity error for each tissue-type computed on sM_{pv} are presented in Table 3. We concluded that less error was introduced in case of GAN_{pv} , for all three tissues. Further, sM_{pv} are more similar to the ground truth in case of GAN_{pv} where smaller shape and intensity errors were introduced. According to Table 3, CSF has a lower DSM than WM and GM for both $GANs$. The rationale behind it is the nature of T1-w images where CSF is difficult to distinguish from the other non-brain tissues.

Evaluation of multi-class voxels. In this experiment, we quantitatively evaluated multi-class voxels, their position and intensity values. Quantitative evaluation was performed by computing DSM between M_{pv} and sM_{pv} for evaluation of their position in sI_{MRI} , while MAE and MSE were computed to measure the intensity error between I_{MRI} and sI_{MRI} . DSM measured in sI_{MRI} generated

from both *GANs* equals the value of one, which implies the location of multi-class voxels is fully preserved in sM_{pv} for both *GANs*. We measured MAE of 0.134 ± 0.017 and MSE of 0.03 ± 0.008 in the multi-class voxels of sI_{MRI}^b . In the case of sI_{MRI}^{pv} , we measured MAE of **0.079 ± 0.024** and MSE of **0.01 ± 0.007** . We also overlaid I_{MRI} with absolute errors, computed voxel-wise, between M_{pv} and sM_{pv} , to provide more information about the localisation and severity of the errors introduced by a *GAN*, segmentation and PV estimation (see Fig. 4). We found that most of the errors happen at tissue boundaries and observed errors of higher value in case of GAN_b . This result illustrates the benefit of using M_{pv} over M_b for the purpose of preserving well defined tissue borders in sI_{MRI} .

According to the presented quantitative results, we obtained **58.96%** smaller MAE and **33.33%** smaller MSE in multi-class voxels of sI_{MRI}^{pv} comparing to sI_{MRI}^b . The presented results support the illustration of absolute errors and strongly suggest that tissue-borders are preserved with higher accuracy in sI_{MRI} generated by GAN_{pv} opposed to GAN_b .

Evaluation of Deep Gray Matter. The region of DGM contains voxels that belong to WM, GM or to both classes. The border between WM and DGM is vaguely defined and hard to segment. Furthermore, in the context of MRI synthesis, a loosely defined or flawed border between WM and DGM makes it easy to distinguish between I_{MRI} and sI_{MRI} . We evaluated the performance of both models in the region of DGM. Quantitative analysis was performed on sI_{MRI} by computing MAE and MSE to measure the error injected by a *GAN*. In the DGM region of sI_{MRI}^b we measured MAE of 0.129 ± 0.021 and MSE of 0.029 ± 0.01 . Yet, in the same ROI of sI_{MRI}^{pv} we measured MAE of **0.108 ± 0.024** and

Table 2. Tissue-wise validation of sI_{MRI} . MAE and MSE are computed between I_{MRI} and sI_{MRI} inside each tissue class.

	MAE			MSE		
	WM	GM	CSF	WM	GM	CSF
GAN_b	0.03 ± 0.003	0.056 ± 0.005	0.046 ± 0.007	0.009 ± 0.001	0.016 ± 0.002	0.028 ± 0.007
GAN_{pv}	0.014 ± 0.003	0.027 ± 0.004	0.032 ± 0.007	0.003 ± 0.001	0.006 ± 0.001	0.022 ± 0.007

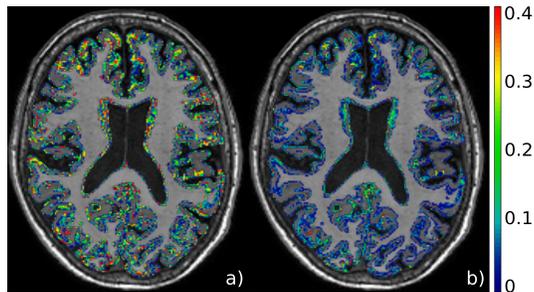
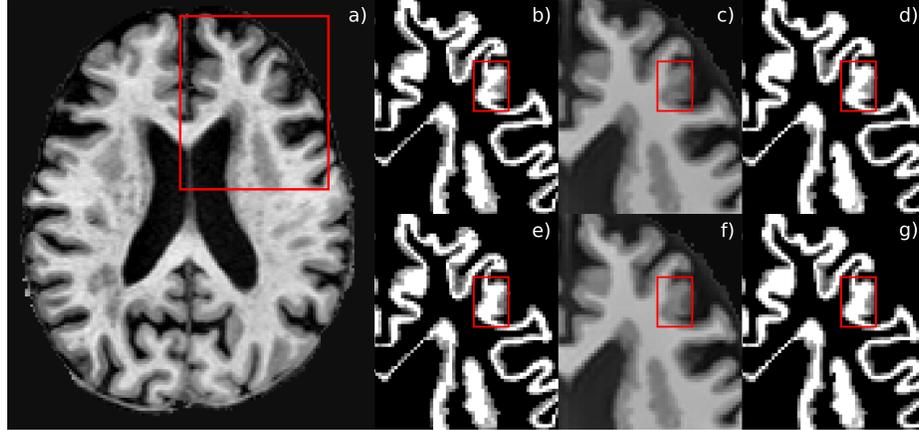


Fig. 4. Location and severity of errors injected into M_{pv} by *GAN*, segmentation and PV estimation. Absolute errors between M_{pv} and sM_{pv}^{Gb} as well as M_{pv} and sM_{pv}^{Gpv} are shown in (a) and (b), respectively.

Table 3. Tissue-wise shape validation of sI_{MRI} and measurements of errors injected into sI_{MRI} by a GAN segmentation and PV estimation.

	DSM			MAE			MSE		
	WM	GM	CSF	WM	GM	CSF	WM	GM	CSF
GAN_b	0.959	0.947	0.922	0.067 ± 0.007	0.098 ± 0.007	0.114 ± 0.018	0.019 ± 0.003	0.028 ± 0.003	0.07 ± 0.018
GAN_{pv}	0.985	0.981	0.954	0.035 ± 0.007	0.047 ± 0.008	0.082 ± 0.017	0.007 ± 0.002	0.009 ± 0.002	0.058 ± 0.017

**Fig. 5.** Introducing fine changes. We introduced a small change into M_{pv} (b) derived from I_{MRI} (a) which is shown in (e). We generated sI_{MRI} from the original and modified M_{pv} shown in (c) and (f). The M_{pv} were then derived from sI_{MRI} in order to verify if the introduced changes are preserved in sI_{MRI} , shown in (d,g).

MSE of 0.022 ± 0.008 . This indicates that the DGM region is more accurately represented in sI_{MRI} generated by the GAN_{pv} when compared to GAN_b .

Introduction of fine changes on PV-map level. The outcomes of this experiment stand for a proof of concept that brain MRI synthesis may be controlled by changing M_{pv} , as the changes are reflected in the sI_{MRI} , while the model still preserves accurate tissue borders. To validate stability, we assessed the ability of the model to preserve fine changes (in this case seven voxels only) in M_{pv} by verifying if the changes are reflected in sI_{MRI} . Both the changed and unchanged M_{pv} were used to generate sI_{MRI} , which were further used to derive sM_{pv} . We obtained the introduced changes in sI_{MRI} and sM_{pv} as shown in Fig. 5.

4 Conclusion

In this work, we tackle the problem of synthesising 3D brain T1-w MRIs with accurate borders between tissues. This is an important feature in the context of medical image applications related to cortical thickness estimation and segmentation. We propose a framework that exploits PV-map properties and demonstrate that it performs better when it comes to synthetic MRI generation with accurate tissue borders compared to binary-map-based alternative. Moreover, we

show that even fine changes introduced on PV-maps are reflected in synthetic images. This implies the possibility of using the framework as a data augmentation mechanism and it will be further explored in our future work.

References

1. Acosta, O., Bourgeat, P., Zuluaga, M.A., Fripp, J., Salvado, O., Ourselin, S., Initiative, A.D.N., et al.: Automated voxel-based 3d cortical thickness measurement in a combined lagrangian–eulerian pde approach using partial volume maps. *Medical image analysis* **13**(5), 730–743 (2009)
2. Bai, W., Oktay, O., Sinclair, M., Suzuki, H., Rajchl, M., Tarroni, G., Glocker, B., King, A., Matthews, P.M., Rueckert, D.: Semi-supervised learning for network-based cardiac mr image segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 253–260. Springer (2017)
3. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: An unsupervised learning model for deformable medical image registration. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 9252–9260 (2018)
4. Ballester, M.Á.G., Zisserman, A.P., Brady, M.: Estimation of the partial volume effect in mri. *Medical image analysis* **6**(4), 389–405 (2002)
5. Frid-Adar, M., Klang, E., Amitai, M., Goldberger, J., Greenspan, H.: Synthetic data augmentation using gan for improved liver lesion classification. In: *2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018)*. pp. 289–293. IEEE (2018)
6. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Advances in neural information processing systems*. pp. 2672–2680 (2014)
7. Han, C., Rundo, L., Araki, R., Nagano, Y., Furukawa, Y., Mauri, G., Nakayama, H., Hayashi, H.: Combining noise-to-image and image-to-image gans: Brain mr image augmentation for tumor detection. *IEEE Access* **7**, 156966–156977 (2019)
8. Hu, Y., Modat, M., Gibson, E., Ghavami, N., Bonmati, E., Moore, C.M., Emberton, M., Noble, J.A., Barratt, D.C., Vercauteren, T.: Label-driven weakly-supervised learning for multimodal deformable image registration. In: *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*. pp. 1070–1074. IEEE (2018)
9. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1125–1134 (2017)
10. Jack Jr, C.R., Bernstein, M.A., Fox, N.C., Thompson, P., Alexander, G., Harvey, D., Borowski, B., Britson, P.J., L. Whitwell, J., Ward, C., 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* **27**(4), 685–691 (2008)
11. Kwon, G., Han, C., Kim, D.s.: Generation of 3d brain mri using auto-encoding generative adversarial networks. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 118–126. Springer (2019)
12. Lee, J.G., Jun, S., Cho, Y.W., Lee, H., Kim, G.B., Seo, J.B., Kim, N.: Deep learning in medical imaging: general overview. *Korean journal of radiology* **18**(4), 570–584 (2017)

13. Lenchik, L., Heacock, L., Weaver, A.A., Boutin, R.D., Cook, T.S., Itri, J., Filippi, C.G., Gullapalli, R.P., Lee, J., Zagurovskaya, M., et al.: Automated segmentation of tissues using ct and mri: a systematic review. *Academic radiology* **26**(12), 1695–1706 (2019)
14. Li, H., Fan, Y.: Non-rigid image registration using self-supervised fully convolutional networks without training data. In: 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018). pp. 1075–1078. IEEE (2018)
15. Mahjourian, R., Wicke, M., Angelova, A.: Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5667–5675 (2018)
16. Mirza, M., Osindero, S.: Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784 (2014)
17. Pombo, G., Gray, R., Varsavsky, T., Ashburner, J., Nachev, P.: Bayesian volumetric autoregressive generative models for better semisupervised learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 429–437. Springer (2019)
18. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention. pp. 234–241. Springer (2015)
19. Ros, G., Sellart, L., Materzynska, J., Vazquez, D., Lopez, A.M.: The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3234–3243 (2016)
20. Roy, A.G., Conjeti, S., Navab, N., Wachinger, C., Initiative, A.D.N., et al.: Quicknat: A fully convolutional network for quick and accurate segmentation of neuroanatomy. *NeuroImage* **186**, 713–727 (2019)
21. Santa Cruz, R., Fernando, B., Cherian, A., Gould, S.: Deeppermnet: Visual permutation learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3949–3957 (2017)
22. Shin, H.C., Tenenholtz, N.A., Rogers, J.K., Schwarz, C.G., Senjem, M.L., Gunter, J.L., Andriole, K.P., Michalski, M.: Medical image synthesis for data augmentation and anonymization using generative adversarial networks. In: International workshop on simulation and synthesis in medical imaging. pp. 1–11. Springer (2018)
23. Taylor, L., Nitschke, G.: Improving deep learning using generic data augmentation. arXiv preprint arXiv:1708.06020 (2017)
24. Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P.: Automated model-based bias field correction of mr images of the brain. *IEEE transactions on medical imaging* **18**(10), 885–896 (1999)
25. Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P.: Automated model-based tissue classification of mr images of the brain. *IEEE transactions on medical imaging* **18**(10), 897–908 (1999)
26. Weiner, M.W., Veitch, D.P., Aisen, P.S., Beckett, L.A., Cairns, N.J., Green, R.C., Harvey, D., Jack Jr, C.R., Jagust, W., Morris, J.C., et al.: The alzheimer’s disease neuroimaging initiative 3: Continued innovation for clinical trial improvement. *Alzheimer’s & Dementia* **13**(5), 561–571 (2017)
27. Xu, Y., Zhu, J.Y., Eric, I., Chang, C., Lai, M., Tu, Z.: Weakly supervised histopathology cancer image segmentation and classification. *Medical image analysis* **18**(3), 591–604 (2014)