



Medical Image Synthesis for Data Augmentation and Anonymization Using Generative Adversarial Networks

Hoo-Chang Shin¹(✉), Neil A. Tenenholtz², Jameson K. Rogers²,
Christopher G. Schwarz³, Matthew L. Senjem³, Jeffrey L. Gunter³,
Katherine P. Andriole², and Mark Michalski²

¹ NVIDIA Corporation, Santa Clara, USA
hshin@nvidia.com

² MGH & BWH Center for Clinical Data Science, Boston, MA, USA
{ntenenholtz,jrogers24}@partners.org

³ Mayo Clinic, Rochester, MN, USA
gunter.jeffrey@mayo.edu

Abstract. Data diversity is critical to success when training deep learning models. Medical imaging data sets are often imbalanced as pathologic findings are generally rare, which introduces significant challenges when training deep learning models. In this work, we propose a method to generate synthetic abnormal MRI images with brain tumors by training a generative adversarial network using two publicly available data sets of brain MRI. We demonstrate two unique benefits that the synthetic images provide. First, we illustrate improved performance on tumor segmentation by leveraging the synthetic images as a form of data augmentation. Second, we demonstrate the value of generative models as an anonymization tool, achieving comparable tumor segmentation results when trained on the synthetic data versus when trained on real subject data. Together, these results offer a potential solution to two of the largest challenges facing machine learning in medical imaging, namely the small incidence of pathological findings, and the restrictions around sharing of patient data.

Keywords: Generative models · GAN · Image synthesis
Deep learning · Brain tumor · Magnetic resonance imaging · MRI
Segmentation

1 Introduction

It is widely known that sufficient data volume is necessary for training a successful machine learning algorithm [6] for medical image analysis. Data with high class imbalance or of insufficient variability [18] leads to poor classification performance. This often proves to be problematic in the field of medical imaging where abnormal findings are by definition uncommon. Moreover, in the case of

image segmentation tasks, the time required to manually annotate volumetric data only exacerbates this disparity; manually segmenting an abnormality in three dimensions can require upwards of fifteen minutes per study making it impractical in a busy radiology practice. The result is a paucity of annotated data and considerable challenges when attempting to train an accurate algorithm. While traditional data augmentation techniques (e.g., crops, translation, rotation) can mitigate some of these issues, they fundamentally produce highly correlated image training data.

In this paper we demonstrate one potential solution to this problem by generating synthetic images using a generative adversarial network (GAN) [9], which provides an additional form of data augmentation and also serves as an effective method of data anonymization. Multi-parametric magnetic resonance images (MRIs) of abnormal brains (with tumor) are generated from segmentation masks of brain anatomy and tumor. This offers an automatable, low-cost source of diverse data that can be used to supplement the training set. For example, we can alter the tumor’s size, change its location, or place a tumor in an otherwise healthy brain, to systematically have the image and the corresponding annotation. Furthermore, GAN trained on a hospital data to generate synthetic images can be used to share the data outside of the institution, to be used as an anonymization tool.

Medical image simulation and synthesis have been studied for a while and are increasingly getting traction in medical imaging community [7]. It is partly due to the exponential growth in data availability, and partly due to the availability of better machine learning models and supporting systems. Twelve recent research on medical image synthesis and simulation were presented in the special issue of Simulation and Synthesis in Medical Imaging [7].

This work falls into the synthesis category, and most related works are those of Chartsias *et al.* [3] and Costa *et al.* [4]. We use the publicly available data set (ADNI and BRATS) to demonstrate multi-parametric MRI image synthesis and Chartsias *et al.* [3] use BRATS and ISLES (Ischemic Stroke Lesion Segmentation (ISLES) 2015 challenge) data set. Nonetheless, evaluation criteria for synthetic images were demonstrated on MSE, SSIM, and PSNR, but not directly on diagnostic quality. Costa *et al.* [4] used GAN to generate synthetic retinal images with labels, but the ability to represent more diverse pathological pattern was limited compared to this work. Also, both previous works were demonstrated on 2D images or slices/views of 3D images, whereas in this work we directly process 3D input/output. The input/output dimension is 4D when it is multi-parametric (T1/T2/T1c/Flair). We believe processing data as 3D/4D in nature better reflects the reality of data and their associated problems.

Reflecting the general trend of the machine learning community, the use of GANs in medical imaging has increased dramatically in the last year. GANs have been used to generate a motion model from a single preoperative MRI [10], upsample a low-resolution fundus image [13], create a synthetic head CT from a brain MRI [16], and synthesizing T2-weight MRI from T1-weighted ones (and vice-versa) [5]. Segmentation using GANs was demonstrated in [21,22].

Finally, Frid-Adar *et al.* leveraged a GAN for data augmentation, in the context of liver lesion classification [8]. To the best of our knowledge, there is no existing literature on the generation of synthetic medical images as form of anonymization and data augmentation for tumor segmentation tasks.

2 Data

2.1 Dataset

We use two publicly available data set of brain MRI:

Alzheimer’s Disease Neuroimaging Initiative (ADNI) Data Set

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 on the ADNI study, see www.adni-info.org. We follow the approach of [17] that is shown to be effective for segmenting the brain atlas of ADNI data. The atlas of white matter, gray matter, and cerebrospinal fluid (CSF) in the ADNI T1-weighted images are generated using the SPM12 [1] segmentation and the ANTs SyN [19] non-linear registration algorithms. In total, there are 3,416 pairs of T1-weighted MRI and their corresponding segmented tissue class images.

Multimodal Brain Tumor Image Segmentation Benchmark (BRATS) Data Set

BRATS utilizes multi-institutional pre-operative MRIs and focuses on the segmentation of intrinsically heterogeneous (in appearance, shape, and histology) brain tumors, namely gliomas [14]. Each patient’s MRI image set includes a variety of series including T1-weighted, T2-weighted, contrast-enhanced T1, and FLAIR, along with a ground-truth voxel-wise annotation of edema, enhancing tumor, and non-enhancing tumor. For more details about the BRATS data set, see braintumorsegmentation.org. While the BRATS challenge is held annually, we used the BRATS 2015 training data set which is publicly available.

2.2 Dataset Split and Pre-processing

As a pre-processing step, we perform skull-stripping [11] on the ADNI data set as skulls are not present in the BRATS data set. The BRATS 2015 training set provides 264 studies, of which we used the first 80% as a training set, and the remaining 20% as a test set to assess final algorithm performance. Hyperparameter optimization was performed within the training set and the test set was evaluated only once for each algorithm and settings assessed. Our GAN operates in 3D, and due to memory and compute constraints, training images were cropped axially to include the central 108 slices, discarding those above and

below this central region, then resampled to $128 \times 128 \times 54$ for model training and inference. For a fair evaluation of the segmentation performance to the BRATS challenge we used the original images with a resolution of $256 \times 256 \times 108$ for evaluation and comparison. However, it is possible that very small tumor may get lost by the downsampling, thus affecting the final segmentation performance.

3 Methods

The image-to-image translation conditional GAN (pix2pix) model introduced in [12] is adopted to translate label-to-MRI (synthetic image generation) and MRI-to-label (image segmentation). For brain segmentation, the generator G is given a T1-weighted image of ADNI as input and is trained to produce a brain mask with white matter, grey matter and CSF. The discriminator D on the other hand, is trained to distinguish “real” labels versus synthetically generated “fake” labels. During the procedure (depicted in Fig. 1(a)) the generator G learns to segment brain labels from a T1-weighted MRI input. Since we did not have an appropriate off-the-shelf segmentation method available for brain anatomy in the BRATS data set, and the ADNI data set does not contain tumor information, we first train the pix2pix model to segment normal brain anatomy from the T1-weighted images of the ADNI data set. We then use this model to perform inference on the T1 series of the BRATS data set. The segmentation of neural anatomy, in combination with tumor segmentations provided by the BRATS data set, provide a complete segmentation of the brain with tumor. The synthetic image generation is trained by reversing the inputs to the generator and training the discriminator to perform the inverse task (i.e., “is this imaging data

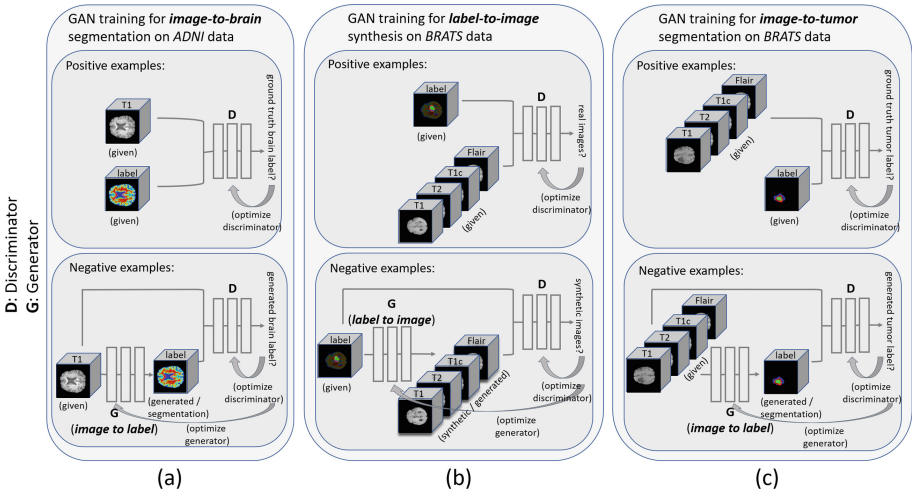


Fig. 1. Illustration of training GAN for (a) MRI-to-brain segmentation; (b) label-to-MRI synthesis; (c) MRI-to-tumor segmentation.

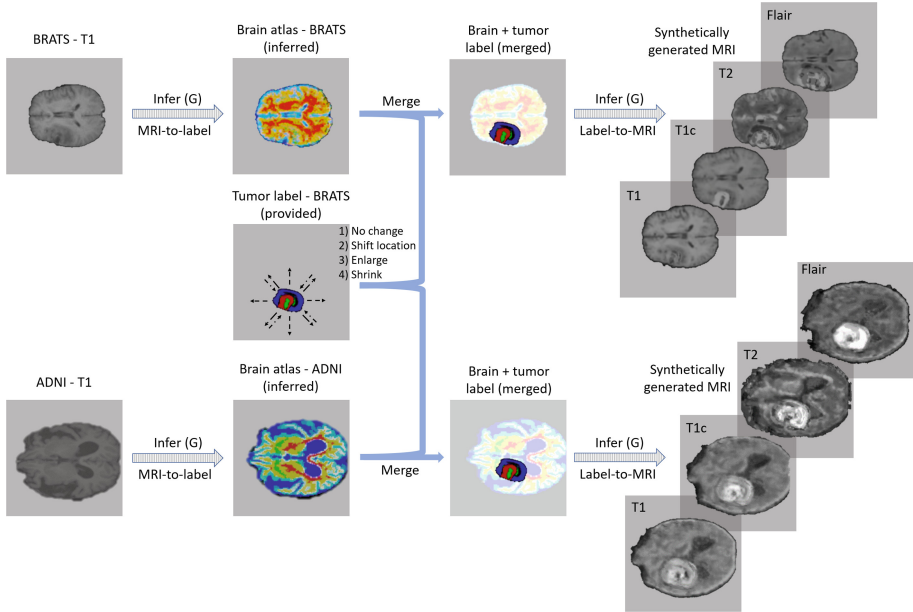


Fig. 2. Workflow of getting synthetic images with variation. On BRATS data set, MRI-to-label image translation GAN is applied to T1-weighted images to get brain atlas. It is then merged with the tumor label given in the BRATS data set, possibly with alterations (shift tumor location; enlarge; shrink). The merged labels (with possibly alterations) are then used as an input to label-to-MRI GAN, to generate synthetic multi-parametric MRI with brain tumor.

acquired from a scanner or synthetically generated?” as opposed to “is this segmentation the ground-truth annotation or synthetically generated?” – Fig. 1(b)). We generate synthetic abnormal brain MRI from the labels and introduce variability by adjusting those labels (e.g., changing tumor size, moving the tumor’s location, or placing tumor on a otherwise tumor-free brain label). Then GAN segmentation module is used once again, to segment tumor from the BRATS data set (input: multi-parametric MRI; output: tumor label). We compare the segmentation performance (1) with and without additional synthetic data, (2) using only the synthetic data and fine-tuning the model on 10% of the real data; and compare their performance of GAN to a top-performing algorithm¹ [20] from the BRATS 2017 challenge.

3.1 Data Augmentation with Synthetic Images

The GAN trained to generate synthetic images from labels allows for the generation of arbitrary multi-series abnormal brain MRIs. Since we have the brain

¹ <https://github.com/taigw/brats17>.

anatomy label and tumor label separately, we can alter either the tumor label or the brain label to get synthetic images with the characteristics we desire. For instance, we can alter the tumor characteristics such as size, location of the existing brain and tumor label set, or place tumor label on an otherwise tumor-free brain label. Examples of this are shown in Fig. 3.

The effect of the brain segmentation algorithm’s performance has not been evaluated in this study.

Since the GAN was first trained on 3,416 pairs of T1-weighted (T1) images from the ADNI data set, generated T1 images are of the high quality, and, qualitatively difficult to distinguish from their original counterparts. BRATS data was used to train the generation of non-T1-weighted image series. Contrast-enhanced T1-weighted images use the same image acquisition scheme as T1-weighted images. Consequently, the synthesized contrast-enhanced T1 images appear reasonably realistic, although higher contrast along the tumor boundary is observed in some of the generated images. T2-weighted (T2) and FLAIR image acquisitions are fundamentally different from the T1-weighted images, resulting in synthetic images that are less challenging to distinguish from scanner-acquired images. However, given a sufficiently large training set on all these modalities, this early evidence suggests that the generation of realistic synthetic images on all the modalities may be possible.

Other than increasing the image resolution and getting more data especially for the sequences other than T1-weighted images, there are still a few important avenues to explore to improve the overall image quality. For instance, more attention likely needs to be paid for the tumor boundaries so it does not look superimposed and discrete when synthetic tumor is placed. Also, performance of brain segmentation algorithm and its ability to generalize across different data sets needs to be examined to obtain higher quality synthetic images combining data sets from different patient population.

The augmentation using synthetic images can be used in addition to the usual data augmentation methods such as random cropping, rotation, translation, or elastic deformation [15]. Moreover, we have more control over the augmented images using the GAN-based synthetic image generation approach, that we have more input-option (i.e., label) to perturb the given image than the usual data augmentation techniques. The usual data augmentation methods rely mostly on random processes and operates on the whole image level than specific to a location, such as tumor. Additionally, since we generate image from the corresponding label, we get more images for training without needing to go through the labor-intensive manual annotation process. Figure 4 shows the process of training GAN with real and synthetic image and label pairs.

3.2 Generating Anonymized Synthetic Images with Variation

Protection of personal health information (PHI) is a critical aspect of working with patient data. Often times concern over dissemination of patient data restricts the data availability to the research community, hindering development of the field. While removing all DICOM metadata and skull-stripping will




Fig. 3. Examples of generated images. The first row depicts the original (“real”) images on which the synthetic tumors were based. Generated images without adjustment of the segmentation label are shown in the second row. Examples of generated images with various adjustments to the tumor segmentation label are shown in the third through fifth rows. The last row depicts examples of synthetic images where a tumor label is placed on a tumor-free brain label from the ADNI data set.

often eliminate nearly all identifiable information, demonstrably proving this to a hospital’s data sharing committee is near impossible. Simply de-identifying the data is insufficient. Furthermore, models themselves are subject to caution when derived from sensitive patient data. It has been shown [2] that private data can be extracted from a trained model.

Development of a GAN that generates synthetic, but realistic, data may address these challenges. The first two rows of Fig. 3 illustrate how, even with the same segmentation mask, notable variations can be observed between the generated and original studies. This indicates that the GAN produces images that do not reflect the underlying patients as individuals, but rather draws individuals from the population in aggregate. It generates new data that cannot be attributed to a single patient but rather an instantiation of the training population conditioned upon the provided segmentation.

4 Experiments and Results

4.1 Data Augmentation Using Synthetic Data

Dice score evaluation of the whole tumor segmentation produced by the GAN-based model and the model of Wang *et al.* [20] (trained on real and synthetic

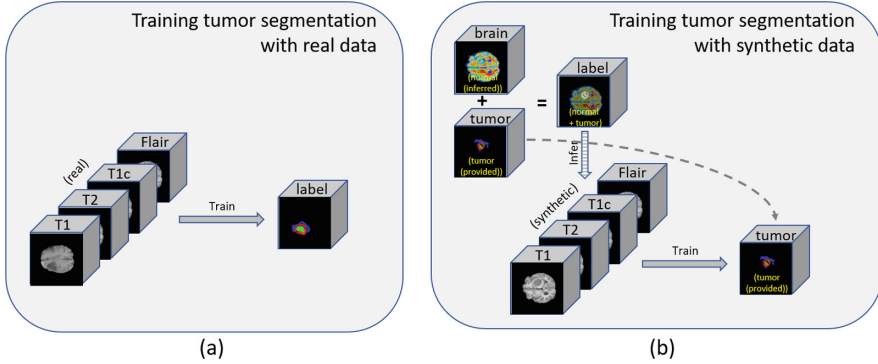


Fig. 4. Training GAN for tumor segmentation with (a) real; and (b) synthetic image-label pairs. Synthetic data generation can increase the training data set with desired characteristics (e.g., tumor size, location, etc.) without the need of labor-intensive manual annotation.

data) are shown in Table 1. The segmentation models are trained on 80% of the BRATS’15 training data only, and the training data supplemented with synthetic data. Dice scores are evaluated on the 20% held-out set from the BRATS’15 training data. All models are trained for 200 epochs on NVIDIA DGX systems.

A much improved performance with the addition of synthetic data is observed without usual data augmentation (crop, rotation, elastic deformation; GAN-based (no-aug)). However, a small increase in performance is observed when added with usual data augmentation (GAN-based (no-aug)), and it applies also to the model of Wang *et al.* [20] that incorporates usual data augmentation techniques.

Wang *et al.* model operates in full resolution (256×256) combining three 2D models for each axial/coronal/sagittal view, whereas our model and generator operates in half the resolution ($128 \times 128 \times 54$) due to GPU memory limit. We up-sampled the GAN-generated images twice the generated resolution for a fair comparison with BRATS challenge, however it is possible that very small tumor may get lost during the down-/up- sampling. A better performance may be observed using the GAN-based model with an availability of GPU with more memory. Also, we believe that the generated synthetic images having half the resolution, coupled with the lack of the image sequences for training other than T1-weighted ones possibly led to the relatively small increase in segmentation performance compared to using the usual data augmentation techniques. We carefully hypothesize that with more T2/Flair images being available, better image quality will be observed for these sequences and so better performance for more models and tumor types.

4.2 Training on Anonymized Synthetic Data

We also evaluated the performance of the GAN-based segmentation on synthetic data only, in amounts greater than or equal to the amount of real data but without including any of the original data. The dice score evaluations are shown in Table 1. Sub-optimal performance is achieved for both our GAN-based and the model of Wang *et al.* [20] when training on an amount of synthetic data equal to the original 80% training set. However, higher performance, comparable to training on real data, is achieved when training the two models using more than five times as much synthetic data (only), and fine-tuning using a 10% random selection of the “real” training data. In this case, the synthetic data provides a form of pre-training, allowing for much less “real” data to be used to achieve a comparable level of performance.

Table 1. Dice score evaluation (mean/standard deviation) of GAN-based segmentation algorithm and BRATS’17 top-performing algorithm [20], trained on “real” data only; real + synthetic data; and training on synthetic data only and fine-tuning the model on 10% of the real data. GAN-based models were trained both with (with aug) and without (no aug) including the usual data augmentation techniques (crop, rotation, translation, and elastic deformation) during training. All models were trained for 200 epochs to convergence.

Method	Real	Real + Synthetic	Synthetic only	Synthetic only, fine-tune on 10% real
GAN-based (no aug)	0.64/0.14	0.80/0.07	0.25/0.14	0.80/0.18
GAN-based (with aug)	0.81/0.13	0.82/0.08	0.44/0.16	0.81/0.09
Wang <i>et al.</i> [20]	0.85/0.15	0.86/0.09	0.66/0.13	0.84/0.15

5 Conclusion

In this paper, we propose a generative algorithm to produce synthetic abnormal brain tumor multi-parametric MRI images from their corresponding segmentation masks using an image-to-image translation GAN. High levels of variation can be introduced when generating such synthetic images by altering the input label map. This results in improvements in segmentation performance across multiple algorithms. Furthermore, these same algorithms can be trained on completely anonymized data sets allowing for sharing of training data. When combined with smaller, institution-specific data sets, modestly sized organizations are provided the opportunity to train successful deep learning models.

References

1. Ashburner, J., Friston, K.J.: Unified segmentation. *Neuroimage* **26**(3), 839–851 (2005)
2. Carlini, N., Liu, C., Kos, J., Erlingsson, Ú., Song, D.: The secret sharer: measuring unintended neural network memorization & extracting secrets. arXiv preprint [arXiv:1802.08232](https://arxiv.org/abs/1802.08232) (2018)
3. Chartsias, A., Joyce, T., Giuffrida, M.V., Tsaftaris, S.A.: Multimodal MR synthesis via modality-invariant latent representation. *IEEE Trans. Med. Imaging* **37**(3), 803–814 (2018)
4. Costa, P.: End-to-end adversarial retinal image synthesis. *IEEE Trans. Med. Imaging* **37**(3), 781–791 (2018)
5. Dar, S.U.H., Yurt, M., Karacan, L., Erdem, A., Erdem, E., Çukur, T.: Image synthesis in multi-contrast MRI with conditional generative adversarial networks. arXiv preprint [arXiv:1802.01221](https://arxiv.org/abs/1802.01221) (2018)
6. Domingos, P.: A few useful things to know about machine learning. *Commun. ACM* **55**(10), 78–87 (2012)
7. Frangi, A.F., Tsaftaris, S.A., Prince, J.L.: Simulation and synthesis in medical imaging. *IEEE Trans. Med. Imaging* **37**(3), 673–679 (2018)
8. Frid-Adar, M., Klang, E., Amitai, M., Goldberger, J., Greenspan, H.: Synthetic data augmentation using GAN for improved liver lesion classification. In: *IEEE International Symposium on Biomedical Imaging (ISBI)* (2018)
9. Goodfellow, I., et al.: Generative adversarial nets. In: *Advances in Neural Information Processing Systems*, pp. 2672–2680 (2014)
10. Hu, Y., et al.: Intraoperative organ motion models with an ensemble of conditional generative adversarial networks. In: Descoteaux, M., et al. (eds.) *MICCAI 2017 Part II. LNCS*, vol. 10434, pp. 368–376. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66185-8_42
11. Iglesias, J.E., Liu, C.-Y., Thompson, P.M., Tu, Z.: Robust brain extraction across datasets and comparison with publicly available methods. *IEEE Trans. Med. Imaging* **30**(9), 1617–1634 (2011)
12. Isola, P., Zhu, J.-Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017
13. Mahapatra, D., Bozorgtabar, B., Hewavitharanage, S., Garnavi, R.: Image super resolution using generative adversarial networks and local saliency maps for retinal image analysis. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) *MICCAI 2017 Part III. LNCS*, vol. 10435, pp. 382–390. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66179-7_44
14. Menze, B.H., et al.: The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE Trans. Med. Imaging* **34**(10), 1993–2024 (2015)
15. Milletari, F., Navab, N., Ahmadi, S.-A.: V-net: fully convolutional neural networks for volumetric medical image segmentation. In: *2016 Fourth International Conference on 3D Vision (3DV)*, pp. 565–571. IEEE (2016)
16. Nie, D., et al.: Medical image synthesis with context-aware generative adversarial networks. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) *MICCAI 2017 Part III. LNCS*, vol. 10435, pp. 417–425. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66179-7_48
17. Christopher, G.S., et al.: A large-scale comparison of cortical thickness and volume methods for measuring Alzheimer’s disease severity. *NeuroImage: Clin.* **11**, 802–812 (2016)

18. Shin, H.-C., et al.: Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning. *IEEE Trans. Med. Imaging* **35**(5), 1285–1298 (2016)
19. Tustison, N.J., et al.: Large-scale evaluation of ANTS and FreeSurfer cortical thickness measurements. *Neuroimage* **99**, 166–179 (2014)
20. Wang, G., Li, W., Ourselin, S., Vercauteren, T.: Automatic brain tumor segmentation using cascaded anisotropic convolutional neural networks. arXiv preprint [arXiv:1709.00382](https://arxiv.org/abs/1709.00382) (2017)
21. Yang, D., et al.: Automatic liver segmentation using an adversarial image-to-image network. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) *MICCAI 2017 Part III*. LNCS, vol. 10435, pp. 507–515. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66179-7_58
22. Zhang, Y., Yang, L., Chen, J., Fredericksen, M., Hughes, D.P., Chen, D.Z.: Deep adversarial networks for biomedical image segmentation utilizing unannotated images. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) *MICCAI 2017 Part III*. LNCS, vol. 10435, pp. 408–416. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66179-7_47