A Preliminary Study on Augmenting Neuroimaging data using a Diffusion Model

Andrea Basile¹, Fabio Calefato¹, Filippo Lanubile¹, Giancarlo Logroscino^{2,3}, Giulio Mallardi^{1,*} and Benedetta Tafuri^{2,3}

¹Dept. of Computer Science, University of Bari Aldo Moro, Bari, Italy ²Dept. of Translational Biomedicine and Neuroscience (DiBraiN), University of Bari Aldo Moro, Bari, Italy ³Center for Neurodegenerative Diseases and the Aging Brain, University of Bari Aldo Moro, Tricase, Italy

Abstract

This preliminary study explores the use of diffusion models for brain imaging generation to address the limitations of small datasets in rare neurodegenerative conditions. Our goal is to improve model robustness by generating realistic variations in medical images. Data scarcity is the main issue for the application of deep learning techniques in neurodegeneration. In the last decade, diffusion models have tried to address this problem as a novel generative technique widely applied for image and video generation. A diffusion model, known for capturing complex data distributions, was trained on a multicenter dataset of Structural Magnetic Resonance Images of healthy subjects to generate a high-quality synthetic dataset. Our results show that the Maximum Mean Discrepancy between two distributions is 0.036, thus indicating that the two distributions are quite similar. However, other metrics such as the Frechet Inception Distance and the Multiscale Structural Similarity Index Measure achieve suboptimal results. Although far from model optimization, these preliminary results demonstrate that diffusion models can be a valid tool to generate high-quality brain imaging data.

Keywords

deep learning, data augmentation, diffusion models, imaging, healthcare

1. Introduction

Magnetic resonance imaging (MRI) represents a valuable technique for investigating brain structure and function. This imaging modality has revolutionized our understanding of various neurological conditions by providing detailed brain images. In the contest of rare neurodegenerative diseases, such as Amyotrophic Lateral Sclerosis (ALS) and Frontotemporal Dementia (FTD), due to the low prevalence of these conditions, the main issue is the lack of a sufficiently high-quality neuroimaging dataset to conduct research studies with robust and generalizable results. This scarcity of data not only impacts the reliability of research in this field but also poses significant challenges for clinical diagnosis and monitoring.

Recent advances in machine learning, particularly the development of diffusion models, offer an effective solution to address these limitations. Diffusion models are capable of capturing complex data distributions and generating high-quality synthetic images. By integrating these models into the neuroimaging workflow, it is possible to augment existing datasets, potentially overcoming the challenges posed by small samples and improving the quality of neuroimaging studies.

This study aims to explore the use of diffusion models to augment neuroimaging data for rare neurodegenerative diseases. We hypothesize that the synthetic images generated by diffusion model will provide high-fidelity MRIs, comparable with real neuroimaging data. Our results could represent a crucial point to significantly improve not only quality and reliability of neuroimaging studies but also the clinical management of conditions like ALS and FTD.

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Vorkshop Proceedings

HC@AIxIA 2024: 3rd AIxIA Workshop on Artificial Intelligence For Healthcare, November 25–28, 2024, Bolzano, Italy *Corresponding author.

Our paper is organized as follows. Section 2 reviews related work on diffusion models for image synthesis. Section 3 presents our proposed method using denoising diffusion probabilistic models (DDPM) to generate synthetic 3D MRI data. Section 4 describes the experimental setup, including datasets and preprocessing steps. Section 5 discusses the results and evaluations of our model. Finally, Section 6 concludes with a summary of the findings and future research directions.

2. Related Work

Diffusion models have emerged as powerful tools for high-quality image synthesis across various domains, surpassing traditional generative models such as GANs. We explore key advancements in the application of diffusion models to general image synthesis, medical imaging, and neuroimaging. Ho et al. [1] introduce diffusion probabilistic models for high-quality image synthesis, demonstrating that these models outperform traditional generative techniques by linking diffusion processes with denoising score matching. Their method achieves impressive results on standard benchmarks such as CIFAR-10. Building on this foundation, Rombach et al. [2] propose Latent Diffusion Models (LDMs), which apply diffusion processes in the latent space of pre-trained autoencoders. This approach significantly reduces computational requirements while maintaining high image quality. It is versatile enough to handle text-to-image synthesis and super-resolution tasks through efficient training and cross-attention mechanisms.

Transitioning to specific applications in the medical domain, Khader et al. [3] propose the integration of denoising diffusion probabilistic models with VQ-GANs to generate realistic and diverse 3D medical images. This approach proves to be effective for data augmentation and enhancing segmentation tasks across various medical imaging datasets, outperforming traditional GAN approaches. Similarly, Dorjsembe et al. [4] present Med-DDPM, a conditional diffusion model designed to generate realistic 3D brain MRI images conditioned on segmentation masks, effectively addressing data scarcity and privacy concerns in medical imaging. Their model not only enhances visual fidelity over existing methods but also improves tumor segmentation accuracy, showcasing its utility for data augmentation and anonymization purposes.

In the context of neurological research, Dhinagar et al. [5] introduce conditional diffusion models tailored to generate synthetic brain MRI images for Alzheimer's disease research. By creating counter-factual images that highlight disease-specific changes, their approach enhances classifier performance and interpretability, supporting the visualization and detection of Alzheimer's effects, which is beneficial for clinical diagnostics and neuroscience studies. Extending this specialized application, Pinaya et al. [6] use LDMs to generate high-resolution synthetic 3D brain MRI images conditioned on demographic and anatomical variables such as age and brain volume. Their method not only achieves superior quality and stability compared to GANs but also produces a large, publicly available dataset of synthetic brain images, facilitating further research in the field.

3. Proposed Approach

We use a denoising diffusion probabilistic model (DDPM) to generate 3D MRI data. Figure 1 shows the proposed model that takes preprocessed T1-weighted MRIs to generate synthetic images, similar to the inputs. The details of the proposed framework are presented in the next section.

3.1. Diffusion for MRI Data Synthesis in Rare Neurodegenerative Diseases

Inspired by non-equilibrium thermodynamics, diffusion models offer a promising alternative to Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs) for high-quality data synthesis. Recent studies, such as those by Dhinagar et al. [5] and Pinaya et al. [6], have effectively applied diffusion models to neuroimaging for Alzheimer's research and high-resolution synthetic brain generation. Building on these advancements, our study adapts the denoising diffusion probabilistic model



Figure 1: Architecture of the denoising diffusion probabilistic model (DDPM) adapted for 3D brain MRI synthesis

(DDPM) to generate realistic 3D brain MRIs tailored specifically for handling limited datasets in rare neurodegenerative conditions.

In contrast to GANs, diffusion models provide better training stability and produce high-fidelity results for audio and graphics [7]. VAEs, instead, tend to have lower image quality than DDPMs, mainly due to their limited expressiveness or imperfect loss criteria. DDPMs, on the other hand, excel in synthesizing more complex and diverse image modalities with less risk of model collapse. However, DDPMs present a unique challenge: they often alter the original data distribution of input images due to introduced random noise, potentially neglecting the structural consistency inherent in the input data. Despite this, DDPMs offer greater flexibility and potential for generating high-quality images than VAEs, especially when capturing complex details and maintaining diversity in the generated images [8].

Ho et al. [1], the authors of the first DDPM, defined a diffusion probabilistic model as a parameterized Markov chain trained using variational inference to produce samples that match the data after a finite time. Transitions of this chain are learned to reverse a diffusion process, which is a Markov chain that gradually adds noise to the data in the opposite direction of sampling until the signal is destroyed. Following the principles outlined in previous studies on diffusion models, we apply conditional Gaussian sampling to preserve key structural elements during the image generation process. This approach provides a stable foundation for generating clinically viable synthetic brain images.

Our approach leverages a customized DDPM model to handle 3D neuroimaging data, introducing specific network architecture and noise schedule adjustments to preserve anatomical accuracy in MRI synthesis. Not previously employed in medical image synthesis, these modifications ensure that the generated images retain structural consistency critical for clinical relevance. In the context of brain MRI synthesis for rare neurological conditions, we begin with a complex distribution of brain structures $p(x_0)$, where x_0 represents a 84×128×84 3D T1-weighted MRI volume. The forward diffusion process, $q(x_1:_T|x_0)$, progressively adds Gaussian noise to the image over *t* timesteps, defined by:

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t x_{t-1}, \beta_t \mathbf{I}})$$
(1)

where β_t is a noise schedule that typically increases with *t*, carefully chosen to balance the trade-off between training stability and sampling speed, this process transforms the intricate distribution of brain images into a simple Gaussian distribution $p(x_T)$, effectively destroying all structural information in the original image.

The core of the DDPM lies in learning the reverse process, $p_{\theta}(x_{t-1}|x_t)$, parameterized by a neural network $\epsilon_{\theta}(x_t, t)$, which estimates the noise added at each step. This network is typically a U-Net

architecture, modified to handle 3D volumetric data and conditioned on the timestep *t*. The training objective is to minimize the difference between the predicted and actual noise:

$$\mathscr{L} = \mathbb{E}_{t,x_0,\epsilon}[\|\epsilon - \epsilon_{\theta}(x_t, t)\|^2]$$
⁽²⁾

where $x_t = \sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$, and $\bar{\alpha}_t = \prod_{s=1}^t (1 - \beta_s)$. This formulation allows for efficient training using a single network to approximate the reverse process for all timesteps.

During inference, we sample $x_T \sim \mathcal{N}(0, \mathbf{I})$ and iteratively apply the learned denoising process:

$$x_{t-1} = \frac{1}{\sqrt{1-\beta_t}} (x_t - \frac{\beta_t}{\sqrt{1-\bar{\alpha}_t}} \epsilon_{\theta}(x_t, t)) + \sigma_t z \tag{3}$$

where $z \sim \mathcal{N}(0, \mathbf{I})$ and σ_t is a slight noise term added to prevent mode collapse.

The application of DDPMs to MRI synthesis for rare brain diseases offers several advantages, particularly in addressing the challenge of small datasets. Rare neurological conditions often result in limited available MRI data, making it challenging to train robust models or conduct comprehensive studies. Several techniques can be employed to overcome this limitation, such as Transfer Learning, Few-shot Learning, Standard data augmentation, and so on.

When applied to DDPMs for rare brain disease MRI synthesis, these techniques may significantly enhance the model's ability to generate diverse, high-quality samples despite limited data availability. DDPMs' progressive denoising process allows for interpretable intermediate results, providing insights into the model's decision-making process—a crucial feature when dealing with rare diseases where every aspect of the generation process needs to be scrutinized.

However, applying DDPMs to 3D MRI synthesis for rare brain diseases also presents unique challenges. The high dimensionality of the volumetric data (84×128×84 in this case) combined with the scarcity of samples requires careful architectural design and regularization strategies to prevent low performance. In addition, ensuring the preservation of disease-specific anatomical details while maintaining overall brain structure consistency requires sophisticated loss functions and potentially incorporating domain-specific knowledge from neurologists and radiologists.

In conclusion, while DDPMs show great promise for MRI data synthesis in the context of rare brain diseases, their successful application requires a careful balance between leveraging their generative power and addressing the specific challenges posed by limited data availability.

3.2. Model Configuration and Training

We implemented a modified version of UNet adapted for 3D data using MONAI¹ [9], an open source framework based on PyTorch and specialized in deep learning in healthcare imaging. Our diffusion model used a 3D UNet architecture with spatial dimensions of 3, 1 input and output channels and a channel structure of [128, 128, 256]. Attention mechanisms were applied at the deepest level, with 256 head channels. The model processed tensors of size $84 \times 128 \times 84$ (depth × height × width) through 2 residual blocks at each level.

We utilized the DDPMScheduler², a class from the MONAI repository that defines the methodology for adding noise to an image, configured as follows:

The "scaled_linear_beta" schedule determines how the noise is added and removed during the diffusion process. This schedule begins by defining a linear schedule for the β_t values, starting from beta_start and ending at beta_end over num_train_timesteps. The linear β_t schedule is then scaled to improve numerical stability, particularly for longer diffusion processes.

¹MONAI: Medical Open Network for AI, https://monai.io/

²DDPMScheduler, https://github.com/Project-MONAI/MONAI/blob/dev/monai/networks/schedulers/ddpm.py

For each timestep *t*, the scheduler computes $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$. During the forward process, noise is gradually added to the image according to these α_t and $\bar{\alpha}_t$ values. In the reverse process (image generation), the scheduler uses these values to guide the gradual denoising of random noise into a coherent brain MRI.

The training process employed a scaled linear beta schedule with 1000 timesteps, beta values ranging from 0.0005 to 0.0195, and 1000 inference steps. We optimized the model using Adam optimizer with a learning rate of 5×10^{-5} over 400 epochs. To accelerate training, we leveraged the 'Accelerate' library for efficient multi-GPU processing.

The training was carried out on LEONARDO, a pre-exascale Tier-0 EuroHPC supercomputer at CINECA³. Specifically, we used the Booster Module partition, which consists of BullSequana X2135 "Da Vinci" GPU Blades. Each node is equipped with a 32-core Intel Xeon Platinum 8358 CPU, 512 GB of RAM, and 4 NVIDIA custom Ampere A100 GPUs with 64GB HBM2e memory each, connected via NVLink 3.0. This configuration allowed us to efficiently train our 3D diffusion model on large-scale medical imaging data, leveraging state-of-the-art hardware and software optimizations.

4. Experimental Setting

4.1. Dataset Details

We initially focused on building a model based on images of healthy subjects to establish its capabilities to capture normal brain structures and their variability. This approach represents the first step for future adaptations to rare neurodegenerative diseases, such as Frontotemporal Dementia (FTD), where datasets are typically much smaller and more specialized.

Through our collaboration with the Center for Neurodegenerative Diseases and the Aging Brain, University of Bari Aldo Moro at Pia Foundation of Cult and Religion "Card. G. Panico" (CMND), we obtained a diverse dataset of T1-weighted healthy brain MRI scans from multiple public sources. This dataset was curated to provide a robust foundation for our initial model development, focusing primarily on healthy subjects. The dataset comprises images from the following sources:

- ADNI (Alzheimer's Disease Neuroimaging Initiative): A longitudinal multicenter study designed to develop clinical, imaging, genetic, and biochemical biomarkers for the early detection and tracking of Alzheimer's disease.
- NIFD (Neuroimaging in Frontotemporal Dementia): A dataset focused on frontotemporal dementia, providing valuable insights into brain structure changes associated with this condition.
- OASIS (Open Access Series of Imaging Studies): A project aimed at making neuroimaging datasets freely available to the scientific community. We utilized data from three OASIS projects:
 - OASIS-1: A cross-sectional collection of young, middle-aged, nondemented, and demented older adults.
 - OASIS-2: A longitudinal collection of older adults with and without dementia.
 - OASIS-3: A compilation of MRI and PET imaging and related clinical data for normal aging and Alzheimer's Disease.
- PPMI (Parkinson's Progression Markers Initiative): A landmark study to identify biomarkers of Parkinson's disease progression.

All datasets are publicly accessible for research and were used in full compliance with ethical guidelines and data privacy standards. All T1-weighted MRI scans from these datasets underwent consistent preprocessing, as detailed in the subsequent subsection, to ensure the uniformity of the images. The final curated dataset consisted of 1,017 preprocessed images. We allocated 80% of these images (approximately 814) for model training, with the remaining 20% (approximately 203) reserved for validation purposes. This split allows for robust model training while retaining a substantial portion of performance evaluation.

³CINECA HPC, https://www.cineca.it/en/

4.2. Preprocessing and Input Representation

To ensure consistency in input orientations and intensities across our datasets, all images underwent a standardized preprocessing pipeline using AssemblyNet [10]. This comprehensive preprocessing involved the following steps [11]:

- 1. **Denoising** Manjón et al. [12]: This step reduced random variations in image intensity, enhancing the signal-to-noise ratio and improving overall image quality.
- 2. **Inhomogeneity correction** Tustison et al. [13]: Also known as bias field correction, this process addressed variations in image intensity caused by magnetic field inhomogeneities, ensuring uniform intensity across the entire brain volume.
- 3. Affine registration to MNI space Avants et al. [14]: Images were spatially aligned to a standard Montreal Neurological Institute (MNI) template. This transformation mapped each brain to a common coordinate system (181 x 217 x 181 voxels at 1 x 1 x 1 mm³ resolution), facilitating inter-subject comparisons.
- 4. Fine Inhomogeneity correction using SPM Ashburner and Friston [15]
- 5. **Tissue-based intensity normalization** Manjón et al. [16]: This step adjusted image intensities based on specific tissue types (e.g., gray matter, white matter), standardizing intensity ranges across different scans and scanners.
- 6. **Brain extraction** Manjón et al. [17]: Non-brain tissues (e.g., skull, scalp) were removed from the images, isolating the brain for subsequent analysis.

Following these steps, we further refined the images by centralizing and normalizing intensities within the brain mask while setting the background to zero. This process ensured that all brain regions were scaled consistently across the dataset.

Finally, to accommodate GPU memory constraints and optimize computational efficiency during model training, we resized the images from their original dimensions (181 x 217 x 181) to 84 x 128 x 84. This adjustment preserved essential structural information while reducing computational load, which is critical given the high memory demands of the diffusion model settings. The memory-intensive nature of these settings limits the model's ability to process full-resolution images on available hardware.

5. Results

In this section, we present the results of our diffusion model.

5.1. Model optimization

In the context of diffusion models, Mean Squared Error (MSE) serves as a crucial loss function during the training process, as it quantifies the difference between the predicted noise and the actual noise added to the data at each timestep of the diffusion process. More specifically:

$$MSE = \frac{1}{n} \sum_{i=1}^{n} (y_i - \hat{y}_i)^2$$
(4)

where y_i represents the actual noise added to the image, and \hat{y}_i is the noise predicted by the model. A lower MSE indicates that the model has become more adept at predicting the noise added during the forward diffusion process, which is essential for effective image generation during the reverse diffusion process. The training of our diffusion model took about six hours and covered 400 epochs. At this point, we observed convergence in our primary loss metric (MSE). The final MSE value achieved was 0,0002, as illustrated in Figure 2.



Figure 2: MSE behavior during the training phase in 400 epochs

5.2. Evaluation metrics

Frechet Inception Distance (FID) [18] calculates the distance between two distributions of feature vectors. This metric was explicitly applied to assess the quality of synthetic images compared to real ones. In order to compute the distance, it is necessary to load a pre-trained model (for example, RadImageNet for 2D and MedicalNet for 3D images), which will extract feature vectors from the images and then compute the statistics like mean and variance used to compute the Frechet distance. A lower value of FID means that the two distributions are similar.

Unbiased Maximum Mean Discrepancy (MMD) [19] is a kernel-based method to measure the similarity between two distributions. It is a non-negative metric where a smaller value indicates a closer match between the two distributions. Multi-Scale Structural Similarity Index Measure (MS-SSIM) [20] is a similarity metric usually used in image generation contexts to measure the structural similarity of data within the same dataset. This index is a value between -1 and 1, where 1 indicates perfect similarity, 0 indicates no similarity, and -1 indicates perfect anti-correlation. We evaluated the metrics over 86 images from both the real images dataset and the synthetic one, achieving the following results:

	MMD	FID	MS-SSIM
medical-3D-DDPM (Ours)	0.036	19.39	0.58
real images	-	-	0.74

Table 1

Comparison of MMD, FID, and MS-SSIM Metrics

The MMD shows promising preliminary results. The value is very close to 0, indicating that the two distributions are pretty similar. However, the FID is higher, suggesting that the features extracted from the real and synthetic datasets are somewhat different. However, the result is promising, given that this is a preliminary study, as depicted in Figure 3.

Lastly, MS-SSIM computed on the synthetic dataset is lower than that of the real dataset, indicating that our model generates sufficiently similar brains. In contrast, the structural similarity in the real dataset is higher, suggesting that the brains within it are approximately 16% more similar to each other than those generated by our model.

Expert neuroradiologists from CMND reviewed a selection of generated images qualitatively, validating their anatomical plausibility and structural consistency, which are crucial for clinical applications. The combination of quantitative metrics and expert validation emphasizes our model's utility and areas for further refinement.





(e) Synthetic Image 2



(d) Synthetic Image 1



(f) Synthetic Image 3

Figure 3: Comparison of real and synthetic MRI shows the diffusion model's anatomical accuracy

6. Conclusion and Future Work

Our work demonstrates the potential of diffusion models in generating synthetic 3D T1-weighted MRI scans of healthy brains. Although the current results are promising, there are several avenues for future research and improvement.

One of the key directions for future research is improving the resolution of the generated MRI images. We acknowledge that the current resolution is insufficient for detecting small-scale neurodegenerative alterations. To achieve this, we plan to implement a latent diffusion model that operates within a compressed latent space. This approach is expected to reduce computational complexity and improve the model's ability to capture both global structures and fine details, leading to higher fidelity in the generated images. Furthermore, this will enable the application of the model to more clinically relevant scenarios where high-resolution imaging is critical.

To broaden our model's applicability to rare neurological conditions, we will employ transfer learning techniques. This approach addresses the challenge of limited data availability for rare brain diseases. We plan to fine-tune our pre-trained model on small, condition-specific datasets and explore few-shot and zero-shot learning methods. Collaborating with clinical partners will ensure the generated images' relevance for rare disorders. This adaptation aims to create valuable tools for research and training in rare neurological conditions.

To integrate our model into healthcare environments, we propose developing an automated Machine Learning Operations (MLOps) pipeline. This pipeline will implement CI/CD practices for medical imaging AI models, including automated data validation, performance monitoring, and rigorous security measures compliant with healthcare regulations. We aim to ensure seamless integration with hospital information systems (HIS). This MLOps pipeline will maintain model accuracy, currency, and deployability in clinical settings, potentially accelerating the adoption of AI-generated synthetic MRI in medical research and practice.

Acknowledgments

This study has been carried out with the co-financing of the Ministry of University and Research in the framework of PNC "DARE - Digital lifelong prevention project" (PNC0000002 - CUP B53C22006450001).

References

- [1] J. Ho, A. Jain, P. Abbeel, Denoising diffusion probabilistic models, in: Proceedings of the 34th International Conference on Neural Information Processing Systems (NeurIPS '20), Curran Associates, Inc., Red Hook, NY, USA, 2020, pp. 6840-6851. URL: https://dl.acm.org/doi/10.5555/ 3495724.3496298.
- [2] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, B. Ommer, High-resolution image synthesis with latent diffusion models, 2022. URL: https://arxiv.org/abs/2112.10752. arXiv:2112.10752.

- [3] F. Khader, G. Müller-Franzes, S. Tayebi Arasteh, T. Han, C. Haarburger, M. Schulze-Hagen, P. Schad, S. Engelhardt, B. Baeßler, S. Foersch, J. Stegmaier, C. Kuhl, S. Nebelung, J. N. Kather, D. Truhn, Denoising diffusion probabilistic models for 3D medical image generation, Scientific Reports 13 (2023) 7303. doi:10.1038/s41598-023-34341-2.
- [4] Z. Dorjsembe, H.-K. Pao, S. Odonchimed, F. Xiao, Conditional diffusion models for semantic 3d brain mri synthesis, IEEE Journal of Biomedical and Health Informatics 28 (2024) 4084–4093. doi:10.1109/JBHI.2024.3385504.
- [5] N. J. Dhinagar, S. I. Thomopoulos, E. Laltoo, P. M. Thompson, Counterfactual mri generation with denoising diffusion models for interpretable alzheimer's disease effect detection, bioRxiv (2024). doi:10.1101/2024.02.05.578983.
- [6] W. H. L. Pinaya, P.-D. Tudosiu, J. Dafflon, P. F. Da Costa, V. Fernandez, P. Nachev, S. Ourselin, M. J. Cardoso, Brain imaging generation with latent diffusion models, in: A. Mukhopadhyay, I. Oksuz, S. Engelhardt, D. Zhu, Y. Yuan (Eds.), Deep Generative Models, Springer Nature Switzerland, Cham, 2022, pp. 117–126. doi:10.1007/978-3-031-18576-2_12.
- [7] P. Dhariwal, A. Nichol, Diffusion models beat gans on image synthesis, 2021. arXiv: 2105.05233.
- [8] S. Dayarathna, K. T. Islam, S. Uribe, G. Yang, M. Hayat, Z. Chen, Deep learning based synthesis of mri, ct and pet: Review and analysis, Medical Image Analysis 92 (2024) 103046. doi:10.1016/j. media.2023.103046.
- [9] M. J. Cardoso, W. Li, R. Brown, N. Ma, E. Kerfoot, Y. Wang, B. Murrey, A. Myronenko, C. Zhao, D. Yang, V. Nath, Y. He, Z. Xu, A. Hatamizadeh, A. Myronenko, W. Zhu, Y. Liu, M. Zheng, Y. Tang, I. Yang, ..., A. Feng, Monai: An open-source framework for deep learning in healthcare, 2022. arXiv:2211.02701.
- [10] P. Coupé, B. Mansencal, M. Clément, R. Giraud, B. Denis de Senneville, V.-T. Ta, V. Lepetit, J. V. Manjon, Assemblynet: A large ensemble of cnns for 3d whole brain mri segmentation, NeuroImage 219 (2020) 117026. doi:10.1016/j.neuroimage.2020.117026.
- [11] J. V. Manjón, P. Coupé, volbrain: An online mri brain volumetry system, Frontiers in Neuroinformatics 10 (2016). doi:10.3389/fninf.2016.00030.
- [12] J. V. Manjón, P. Coupé, L. Martí-Bonmatí, D. L. Collins, M. Robles, Adaptive non-local means denoising of mr images with spatially varying noise levels, Journal of Magnetic Resonance Imaging 31 (2010) 192–203.
- [13] N. J. Tustison, B. B. Avants, P. A. Cook, Y. Zheng, A. Egan, P. A. Yushkevich, J. C. Gee, N4itk: Improved n3 bias correction, IEEE Transactions on Medical Imaging 29 (2010) 1310–1320. doi:10. 1109/TMI.2010.2046908.
- [14] B. B. Avants, N. J. Tustison, G. Song, P. A. Cook, A. Klein, J. C. Gee, A reproducible evaluation of ants similarity metric performance in brain image registration, NeuroImage 54 (2011) 2033–2044. doi:10.1016/j.neuroimage.2010.09.025.
- [15] J. Ashburner, K. J. Friston, Unified segmentation, neuroimage 26 (2005) 839-851.
- [16] J. V. Manjón, J. Tohka, G. García-Martí, J. Carbonell-Caballero, J. J. Lull, L. Martí-Bonmatí, M. Robles, Robust mri brain tissue parameter estimation by multistage outlier rejection, Magnetic Resonance in Medicine 59 (2008) 866–873. doi:10.1002/mrm.21521.
- [17] J. V. Manjón, S. F. Eskildsen, P. Coupé, J. E. Romero, D. L. Collins, M. Robles, Nonlocal intracranial cavity extraction, International journal of biomedical imaging 2014 (2014) 820205.
- [18] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, S. Hochreiter, Gans trained by a two timescale update rule converge to a local nash equilibrium, in: Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17, Curran Associates Inc., Red Hook, NY, USA, 2017, p. 6629–6640.
- [19] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, A. Smola, A kernel two-sample test, J. Mach. Learn. Res. 13 (2012) 723–773.
- [20] Z. Wang, E. Simoncelli, A. Bovik, Multiscale structural similarity for image quality assessment, in: The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003, volume 2, 2003, pp. 1398–1402 Vol.2. doi:10.1109/ACSSC.2003.1292216.