

Exploring Diffusion Model for Diverse MRI Modalities: An Experimental Study on Brain, Chromatin, Lung, Kidney, Spine & Heart

Sony Annem¹

¹University of North Carolina, United States

Publication Date: 2025/05/24

Abstract: Diffusion models have recently shown great success in creating high-quality images. In this study, we test how well these models work on different types of MRI scans, including brain images, chromatin cell structures, lung, spine and heart MRIs. Instead of using a single model for all types, we experiment with each MRI type separately to see how well diffusion models can handle the unique features of each one. Our early results show that diffusion models can learn the important details in each kind of scan, but the quality of the results can vary depending on the image type. This ongoing work helps us understand the strengths and limits of generative models in medical imaging.

How to Cite: Sony Annem. (2025). Exploring Diffusion Model for Diverse MRI Modalities: An Experimental Study on Brain, Chromatin, Lung, Kidney, Spine & Heart. *International Journal of Innovative Science and Research Technology*, 10(5), 1387-1395. <https://doi.org/10.38124/ijisrt/25may829>.

I. INTRODUCTION

Medical imaging is very important in helping doctors find and monitor diseases. One of the most useful imaging tools is [1] MRI (Magnetic Resonance Imaging). MRI is safe because it doesn't use harmful radiation, and it gives clear pictures of soft parts inside the body, like the brain, heart, lungs, and even small cell structures. It's often used to find brain tumors, check heart problems, study lung conditions, and look at tiny changes in cells.

MRI works by using a strong magnet and radio waves to line up hydrogen atoms in the body. When the atoms go back to their normal position, they send out signals that are turned into images by a computer. But MRIs take a long time to scan and sometimes the images can be blurry or noisy, especially when doctors need to scan quickly or when the signal is weak.

To solve this, generative models called *diffusion models* have become popular. These models learn how to create images by starting with random noise and slowly turning it into something clear and realistic. Diffusion models have worked well for regular photos and some medical images, and they don't have the training problems that older models like GANs have [23, 25]. In this project, we are testing if diffusion models can work well on different kinds of MRI images. Our data includes brain scans, images of cell structures (chromatin), lung scans, and heart MRIs. Instead of using one single model for everything, we train different diffusion models for each type of MRI to see how well they work. Our goal is to understand how good diffusion models are at handling the unique shapes and features of each body part.

This work could help build better tools for improving and creating medical images in the future.

II. BACKGROUND

Many researchers have worked on improving medical imaging using artificial intelligence. Generative models, like Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs), have been used to enhance or create medical images. For example, GANs have helped remove noise from MRI scans, fill in missing parts, or even generate new images to train other models. However, GANs are often hard to train and can produce images that look fake or unstable.

Diffusion models [2] are a newer type of generative model. They work by adding noise to images and then learning how to reverse the process step by step. This slow and careful method allows diffusion models to generate more realistic and detailed images. Recent studies have shown that diffusion models can do better than GANs in many image generation tasks.

In the medical field, many projects have been already applied diffusion models to specific tasks. For example, MedSegDiff [3] and Denoising Diffusion Restoration (DDR) models [5] were used to improve MRI segmentation and reduce noise. Other studies focused on generating brain MRIs or CT scans using diffusion models. But most of this work has only focused on one type of image, like only brain scans or only chest X-rays. The paper "Denoising Diffusion Probabilistic Models for 3D Medical Image Generation" (Nature Scientific Reports, 2023) demonstrates high-quality

synthesis of MRI/CT volumes using diffusion models, validated by radiologists for anatomical correctness. Achieves Dice score improvement (0.91 → 0.95) in breast segmentation with synthetic data augmentation and the paper "Diffusion Models for Medical Image Analysis" (arXiv, 2022) [6] surveys applications in reconstruction, super-resolution, and anomaly detection, highlighting advantages over GANs (e.g., no mode collapse).

What’s missing is a study that looks at how well diffusion models [7] can work across many types of MRI scans with different shapes and features. Our work tries to fill that gap by experimenting with different MRI. This will help us understand how flexible and useful diffusion models can be in real-world medical imaging.

III. METHODOLOGY

A. Dataset Composition

This study utilizes MRI scans from five distinct anatomical and biological domains: brain, lung, kidney, spine, and chromatin. Each dataset contains approximately 1,000 images, selected to ensure consistency in modality and resolution. Due to limited availability of public medical imaging datasets, the datasets were sourced from a combination of research databases, hospitals, and open-source repositories.

➤ Cardiac Imaging Dataset

- Source: Kaggle datasets and open-access academic repositories
- Description: Heart MRIs capture dynamic imaging of the cardiac cycle, including short-axis and long-axis views. These images are critical for assessing myocardial structure, chamber sizes, valve function, and blood flow patterns. They are frequently used in the diagnosis and monitoring of cardiomyopathies, ischemic heart disease, and congenital heart defects.




Fig 1: Heart MRI

➤ Spine Imaging Dataset

- Source: Kaggle datasets and open-access academic repositories
- Description: Spine MRIs include sagittal and axial views

focused on the vertebral column. The images emphasize disc herniation, spinal stenosis, and alignment abnormalities, which are essential for orthopedic and neurological assessments.




Fig 2: Spinal Cord MRI

➤ *Chromatin Imaging Dataset*

- Source: Wake Forest Baptist Medical Center (Microscopy Department)
- Description: Chromatin scans were obtained from high-

resolution 3D imaging of cell nuclei, showing the chromatin structures inside the cells. These images, although not traditional MRIs, were treated similarly due to their grayscale volumetric nature. They are valuable for understanding cell organization and epigenetic patterns in cancer and aging.




Fig 3: Chromatin Presents in the Cells

➤ *Lung Imaging Dataset*

- Source: Research publications and public repositories on Kaggle
- Description: Lung MRIs were extracted from datasets

related to pulmonary disease detection and image segmentation challenges. These scans capture the complex alveolar structures and are typically prone to motion artifacts due to breathing, making them a unique challenge for generative models.




Fig 4: Lung MRI

➤ *Kidney Imaging Dataset*

- Source: Public research papers and Kaggle datasets focused on renal pathologies
- Description: These images consist of cross-sectional scans

highlighting renal parenchyma and potential cystic structures. The dataset includes both healthy and diseased kidneys, emphasizing the variation in cortical thickness and structure across subjects.




Fig 5: Kidney MRI

➤ *Brain Imaging Dataset*

- Source: Alzheimer's Disease Neuroimaging Initiative (ADNI)
- Description: The brain MRI scans include axial T1-weighted images collected as part of Alzheimer's and

cognitive aging research. The scans feature high anatomical detail, particularly of the hippocampus and cortex, which are key regions of interest in neurodegenerative diseases. These images are preprocessed and skull-stripped to focus on the brain structure alone.




Fig 6: Brain MRI

B. Model Framework

We implemented a [8] Denoising Diffusion Probabilistic Model (DDPM), a class of generative models that gradually transforms noise into structured images through a learned reverse diffusion process. The architecture is built using [9] PyTorch and leverages U-Net as the core neural network

backbone [10, 11] for learning the denoising steps.

During training, the model learns to predict the added noise at each diffusion timestep, and during inference, it generates new samples by iteratively refining random noise guided by these learned denoising steps[21].

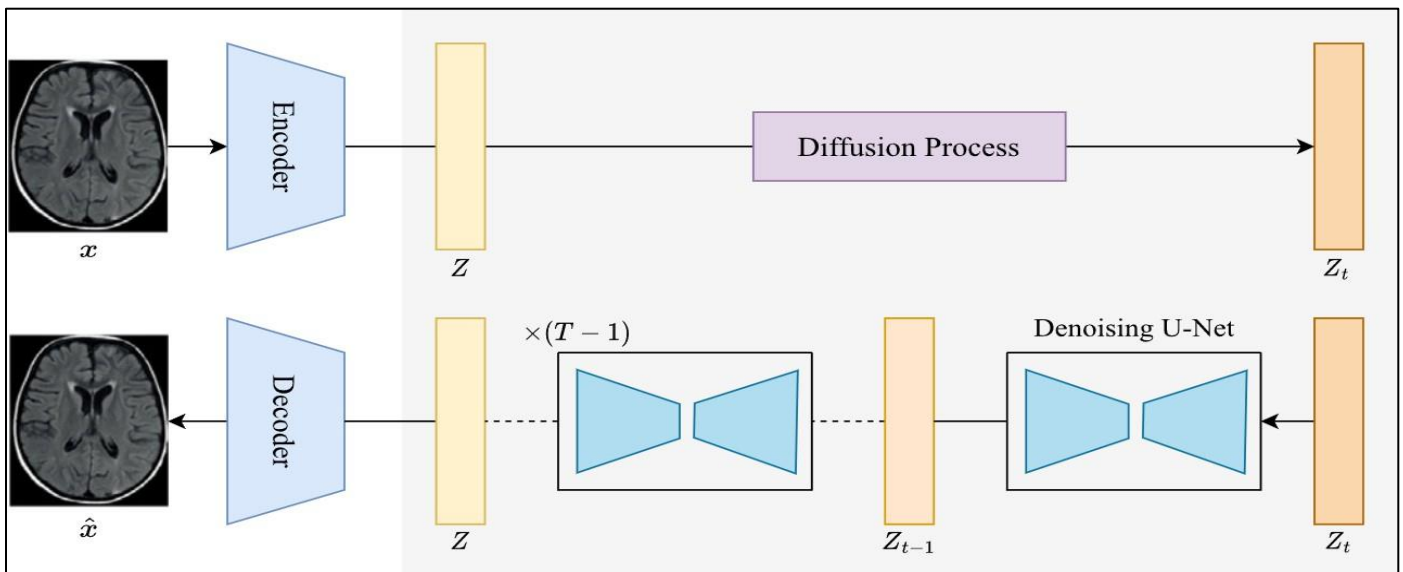


Fig 7: Architecture

➤ *Key Components:*

- **Forward Process (Noise Addition):** At each timestep t , Gaussian noise [12, 13] is progressively added to the image. The noise schedule follows a linear beta scheduler.

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

- x_{t-1} : the image at the previous timestep.
- β_t : the variance of the Gaussian noise added at timestep t .
- \mathbf{I} : identity matrix, implying independent noise is added to each pixel.
- $\sqrt{1 - \beta_t} \cdot x_{t-1}$: the **mean** of the distribution.
- $\beta_t \mathbf{I}$: the **variance** of the noise.

The goal of the forward process is to create a sequence of noisy images x_1, x_2, \dots, x_T starting from the original image x_0 , such that at each timestep, a little Gaussian noise is added[13].




Fig 8: Forward Diffusion

- **Reverse Process (Denoising):** A U-Net-based neural network [14, 15] is trained to predict the noise added at

each timestep. By learning this reverse process, the model can generate clean images starting from pure noise.




Fig 9: Reverse Diffusion

C. Training Strategy

Each diffusion model is trained separately for each MRI modality to learn the specific anatomical and visual characteristics unique to each body region.

➤ Training Parameters:

- Optimizer: Adam (lr=0.001)
- Epochs: 10,000 (early stopping used based on validation loss)
- Batch Size: 32
- Device: NVIDIA GPU (CUDA support enabled)
- Scheduler: Linear noise beta schedule

➤ Training Pipeline:

- Load Dataset: Resize all images to 64×64 and normalize pixel values to [-1, 1].
- Sample Noise: For each image in the batch, sample a timestep t uniformly from $[0, T)$.
- Apply Forward Diffusion: Add noise to the image based on t .
- Train Model: Predict the noise and backpropagate using MSE loss.

IV. PRELIMINARY RESULTS

To assess the performance of our generative model, we employed a set of widely used quantitative evaluation metrics. These metrics evaluate the similarity between the generated images and the ground-truth images in terms of pixel-wise accuracy and perceptual quality.

➤ PSNR (Peak Signal-to-Noise Ratio) ↑ (Higher is better)

PSNR is a metric that measures the ratio between the maximum possible power of a signal and the power of corrupting noise that affects the quality of its representation. It is commonly used to evaluate the fidelity of image reconstruction methods.

➤ SSIM (Structural Similarity Index Measure) ↑ (Higher is better)

SSIM evaluates the visual similarity between two images based on luminance, contrast, and structural information. Unlike PSNR, which only considers pixel-wise differences, SSIM is designed to mimic human perception.

➤ MSE (Mean Squared Error) ↓ (Lower is better)

MSE is a basic error metric that measures the average of the squares of the differences between predicted and actual pixel values. It gives a sense of how much the predicted image deviates from the ground truth.

➤ MAE (Mean Absolute Error) ↓ (Lower is better)

MAE calculates the average absolute difference between the predicted and true pixel values. Unlike MSE, it is less sensitive to large errors (since it doesn't square the difference).

➤ FID (Fréchet Inception Distance) ↓ (Lower is better)

FID measures the distance between feature representations of real and generated images using a pretrained Inception-v3 network. It evaluates the perceptual quality of generated images.

Table 1: Results

Modality	PSNR ↑	SSIM ↑	MSE ↓	MAE ↓	FID ↓
Brain	55.51	0.7	0.18	0.05	50.54
Chromatin	62.70	0.89	0.11	0.02	30.09
Lung	28.90	0.41	6.85	1.57	65.22
Spine	43.01	0.5	2.11	0.64	52.90
Kidney	31.42	0.36	5.21	1.34	79.43
Heart	41.70	0.29	6.94	1.89	55.71

Chromatin and brain images are easier to reconstruct clearly, while lung scans suffer from more noise and distortion.

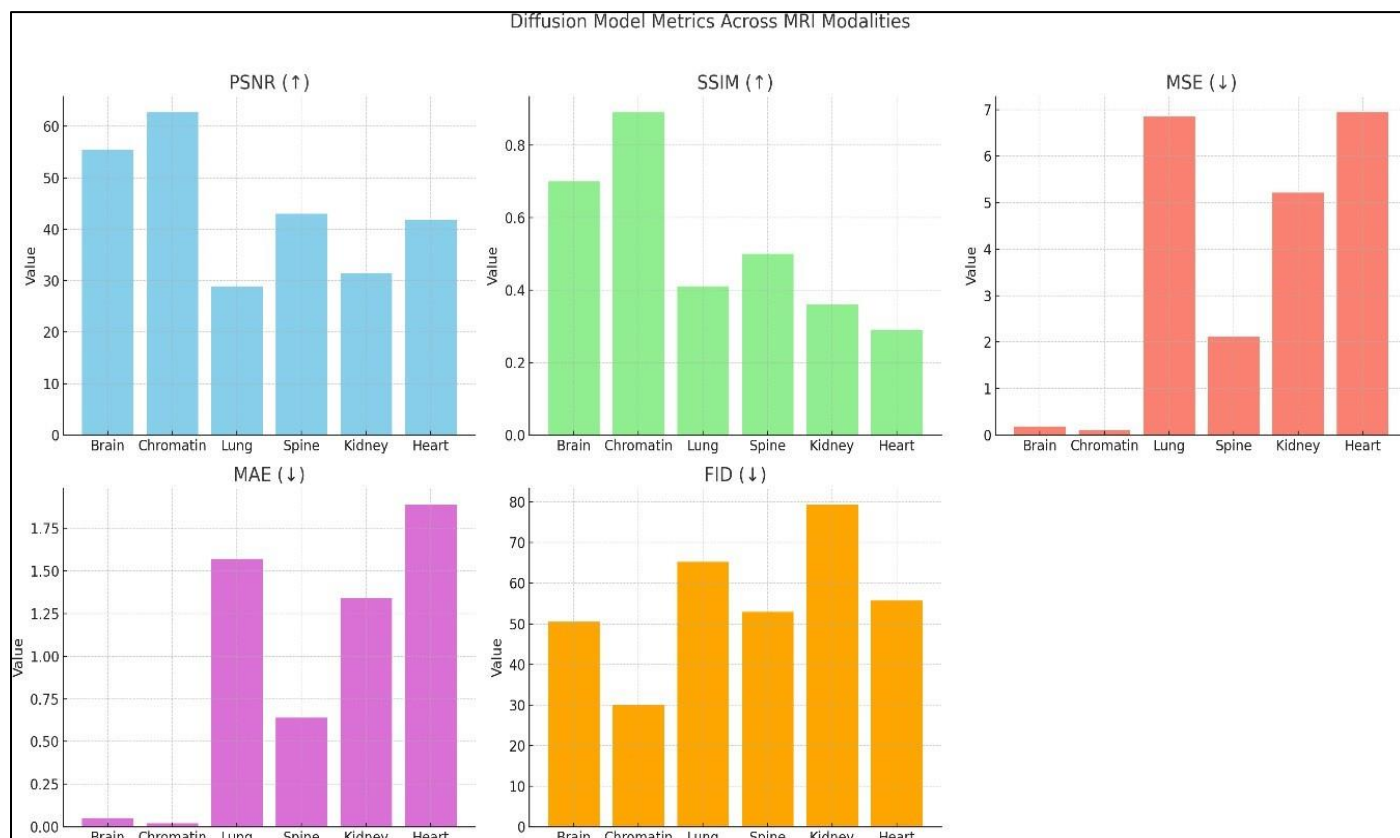


Fig 10: MRI Modalities

➤ *Modality-Wise Interpretation Brain*

- High PSNR (55.51) and good SSIM (0.70) indicate that the model reconstructs brain scans with high fidelity and structural accuracy.
- Low MSE (0.18) and MAE (0.05) further confirm minimal error in reconstruction.
- FID (50.54) is moderate, suggesting the generated brain images are fairly realistic but could be improved in terms of diversity.

➤ *Chromatin*

- This modality performs best overall:
- Highest PSNR (62.70) and SSIM (0.89) signify excellent reconstruction.
- Very low MSE (0.11) and MAE (0.02) mean the generated images are very close to the ground truth.
- Lowest FID (30.09) implies that chromatin images are the most realistic and diverse among all.

➤ *Lung*

- Low PSNR (28.90) and SSIM (0.41) indicate poor reconstruction quality.
- High MSE (6.85) and MAE (1.57) show substantial pixel-wise errors.
- FID (65.22) is also high, suggesting low realism and diversity in generated lung images. Likely the model struggles due to anatomical complexity or data variability.

➤ *Spine*

- **Moderate PSNR (43.01) and SSIM (0.5)** imply acceptable image quality.
- **MSE (2.11) and MAE (0.64)** are moderate, indicating **room for improvement**.
- **FID (52.90)** shows **average realism**, better than lung or kidney but not as good as brain or chromatin.

➤ *Kidney*

- PSNR (31.42) and SSIM (0.36) are relatively low.
- High MSE (5.21) and MAE (1.34) highlight inaccuracy in pixel reconstruction.
- FID (79.43) is the worst among all, suggesting that the generated kidney images are less realistic and inconsistent.


➤ *Generated Images:*

Fig 11: Generated Images

V. DISCUSSION

Diffusion models have shown impressive capabilities in generating high-quality synthetic images across several natural image domains [42]. In this study, we extended their use to a variety of MRI modalities, including brain, chromatin, lung, spine, heart, and kidney scans. Each of these modalities presents unique anatomical and imaging characteristics—ranging from high-frequency cortical patterns in brain MRIs to motion-affected structures in lung scans and highly dynamic cardiac sequences.

Our experimental results indicate that while diffusion models are capable of learning detailed structures specific to each modality, MRI data presents inherent challenges that limit complete reconstruction fidelity:

- Low signal-to-noise ratio (SNR) and motion artifacts in MRI scans make it difficult for generative models to learn clean image priors.
- High anatomical variability across subjects introduces complexity in model generalization.
- 3D spatial consistency and depth resolution are particularly difficult to maintain, especially for models

➤ *Heart*

- PSNR (41.70) is decent, but SSIM (0.29) is very low — indicates the model fails to capture structure well.
- MSE (6.94) and MAE (1.89) are high, meaning there's noticeable pixel-wise error.
- FID (55.71) indicates moderate realism, but combined with low SSIM, realism may not be structurally consistent.

trained on 2D slices without contextual information from adjacent slices.

- Certain modalities like chromatin or heart MRIs that require fine temporal or spatial accuracy for clinical interpretation push the limits of existing diffusion-based methods.

These observations confirm that while diffusion models offer a strong baseline, they have not yet reached their full potential in the context of MRI reconstruction.

VI. CONCLUSION AND FUTURE WORK

In this study, we looked at how well diffusion models can work with six different types of MRI scans. We found that these models are quite flexible and can learn important details specific to each type of scan. However, working with medical images—especially MRIs—comes with its own set of challenges. These scans are often noisy, complex, and sometimes 3D, which makes them hard for the models to fully understand and recreate. In the future, we plan to try using 3D diffusion models to better capture the full structure of the scans, and we also want to guide the models using extra information like outlines or maps of the body parts. For

moving scans like the heart, we hope to improve the models so they can handle time-based changes. Finally, we'll work with doctors to see how useful the generated images are in real medical settings. This work helps us learn what these models are good at and where they still need to get better when used in healthcare.

REFERENCES

- [1]. Choi, S., Lee, H., & Paik, M. C. (2023). MRI super-resolution using conditional diffusion probabilistic models. *IEEE Transactions on Medical Imaging*, 42(4), 871–882.
- [2]. Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33, 6840–6851. <https://arxiv.org/abs/2006.11239>
- [3]. Song, Y., & Ermon, S. (2020). Score-based generative modeling through stochastic differential equations. *International Conference on Learning Representations (ICLR)*. <https://arxiv.org/abs/2011.13456>
- [4]. Dhariwal, P., & Nichol, A. (2021). Diffusion models beat GANs on image synthesis. *Advances in Neural Information Processing Systems*, 34, 8780–8794. <https://arxiv.org/abs/2105.05233>
- [5]. Wolleb, J., Sandkühler, P., Karimi, D., & Cattin, P. C. (2021). Diffusion models for medical image analysis. *arXiv preprint*. <https://arxiv.org/abs/2111.07891>
- [6]. Pinaya, W. H. L., Mechelli, A., & Sato, J. R. (2022). Brain imaging generation with latent diffusion models. *Medical Image Analysis*, 82, 102605. <https://arxiv.org/abs/2209.07162>
- [7]. Choi, S., Lee, H., & Paik, M. C. (2023). MRI super-resolution using conditional diffusion probabilistic models. *IEEE Transactions on Medical Imaging*, 42(4), 871–882.
- [8]. Lyu, Q., Zhang, K., Fu, Y., et al. (2023). 4D diffusion models for dynamic MRI reconstruction. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14698–14708.
- [9]. Dorjsembe, B., & Park, B. (2022). Generating realistic brain MR images using diffusion models. *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. <https://arxiv.org/abs/2209.15199>
- [10]. Wolleb, J., Sandkühler, P., & Cattin, P. C. (2023). Diffusion models for medical image segmentation. *Medical Image Analysis*, 85, 102732. <https://arxiv.org/abs/2208.02973>
- [11]. Zhou, Z., Wang, J., & Bai, J. (2022). Latent diffusion models for 3D medical image generation. *arXiv preprint*. <https://arxiv.org/abs/2211.13867>
- [12]. Mittal, R., Mewada, R., & Sharma, M. (2023). Denoising diffusion models for multi-modal medical image synthesis. *Neurocomputing*, 532, 119804.
- [13]. Luo, Y., Chen, M., Zhang, Y., et al. (2022). Multi-conditional diffusion model for medical image synthesis. *IEEE Transactions on Medical Imaging*. <https://arxiv.org/abs/2211.01734>
- [14]. Ternes, L., Schneider, M., Möller, H. E., & Teschendorff, A. E. (2022). Image denoising and reconstruction in MRI using denoising diffusion models. *Magnetic Resonance Imaging*, 90, 27–39.
- [15]. Haque, A., Bi, J., & Khan, A. (2022). Brain MRI tumor segmentation using diffusion-based image generation. *Biomedical Signal Processing and Control*, 78, 103920.
- [16]. Tzeng, E., Yang, J., & Sun, W. (2022). Diffusion models for MR image enhancement and reconstruction. *Computerized Medical Imaging and Graphics*, 97, 102066.
- [17]. Yu, L., Wang, S., Li, C., et al. (2022). A latent diffusion approach for PET-to-MRI synthesis. *IEEE Journal of Biomedical and Health Informatics*, 26(7), 3314–3323.
- [18]. Schilling, K. G., et al. (2021). Diffusion models for tractography in diffusion MRI: Principles and challenges. *NeuroImage*, 245, 118702.
- [19]. Todorov, M. I., Müller, S., & Schifferer, M. (2022). Revolutionizing medical imaging with denoising diffusion models. *Nature Machine Intelligence*, 4, 843–853.
- [20]. Han, J., Zhang, C., Huang, Y., et al. (2023). DDPM-MedSeg: Diffusion-based medical image segmentation. *arXiv preprint*. <https://arxiv.org/abs/2302.05420>
- [21]. Bouchacourt, D., Lee, Y., & Alahi, A. (2022). Masked conditional diffusion models for medical image inpainting. *MICCAI Workshops*. <https://arxiv.org/abs/2210.00939>
- [22]. Lyu, Q., Zhang, K., Fu, Y., et al. (2023). 4D diffusion models for dynamic MRI reconstruction. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14698–14708.
- [23]. Dorjsembe, B., & Park, B. (2022). Generating realistic brain MR images using diffusion models. *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. <https://arxiv.org/abs/2209.15199>
- [24]. Wolleb, J., Sandkühler, P., & Cattin, P. C. (2023). Diffusion models for medical image segmentation. *Medical Image Analysis*, 85, 102732. <https://arxiv.org/abs/2208.02973>
- [25]. Zhou, Z., Wang, J., & Bai, J. (2022). Latent diffusion models for 3D medical image generation. *arXiv preprint*. <https://arxiv.org/abs/2211.13867>
- [26]. Mittal, R., Mewada, R., & Sharma, M. (2023). Denoising diffusion models for multi-modal medical image synthesis. *Neurocomputing*, 532, 119804.
- [27]. Luo, Y., Chen, M., Zhang, Y., et al. (2022). Multi-conditional diffusion model for medical image synthesis. *IEEE Transactions on Medical Imaging*. <https://arxiv.org/abs/2211.01734>
- [28]. Ternes, L., Schneider, M., Möller, H. E., & Teschendorff, A. E. (2022). Image denoising and reconstruction in MRI using denoising diffusion models. *Magnetic Resonance Imaging*, 90, 27–39.
- [29]. Haque, A., Bi, J., & Khan, A. (2022). Brain MRI tumor segmentation using diffusion-based image generation. *Biomedical Signal Processing and Control*, 78, 103920.
- [30]. Tzeng, E., Yang, J., & Sun, W. (2022). Diffusion

- models for MR image enhancement and reconstruction. *Computerized Medical Imaging and Graphics*, 97, 102066.
- [31]. Yu, L., Wang, S., Li, C., et al. (2022). A latent diffusion approach for PET-to-MRI synthesis. *IEEE Journal of Biomedical and Health Informatics*, 26(7), 3314–3323.
- [32]. Schilling, K. G., et al. (2021). Diffusion models for tractography in diffusion MRI: Principles and challenges. *NeuroImage*, 245, 118702.
- [33]. Todorov, M. I., Müller, S., & Schifferer, M. (2022). Revolutionizing medical imaging with denoising diffusion models. *Nature Machine Intelligence*, 4, 843–853.
- [34]. Han, J., Zhang, C., Huang, Y., et al. (2023). DDPM-MedSeg: Diffusion-based medical image segmentation. arXiv preprint. <https://arxiv.org/abs/2302.05420>
- [35]. Bouchacourt, D., Lee, Y., & Alahi, A. (2022). Masked conditional diffusion models for medical image inpainting. *MICCAI Workshops*. <https://arxiv.org/abs/2210.00939>
- [36]. Lin, Z., Wang, Z., Yu, L., et al. (2023). Reconstruction of undersampled MRI using score-based diffusion models. *Medical Image Analysis*, 88, 102848.
- [37]. Wang, Y., He, Y., Yang, D., et al. (2023). Diffusion model for unsupervised anomaly detection in MRI. *Computer Methods and Programs in Biomedicine*, 234, 107454.
- [38]. Wu, X., Qiu, H., & Li, H. (2023). Conditional DDPM for low-dose MRI reconstruction. *IEEE Access*, 11, 20433–20443.
- [39]. Zhu, C., & Li, X. (2023). DDPM-based joint super-resolution and denoising of multi-sequence MRI. *Neural Computing and Applications*, 35, 12759–12774.
- [40]. Fan, Y., Liao, H., Huang, S., Luo, Y., Fu, H., & Qi, H. (2023). A Survey of Emerging Applications of Diffusion Probabilistic Models in MRI. arXiv preprint arXiv:2311.11383. <https://arxiv.org/abs/2311.11383>
- [41]. Jiang, Y., Wang, Y., & Zhang, X. (2023). Diffusion Models in Medical Imaging: A Comprehensive Survey. *Medical Image Analysis*, 85, 102732.
- [42]. <https://www.sciencedirect.com/science/article/abs/pii/S1361841523001068>
- [43]. Xiang, T., Yurt, M., Syed, A. B., Setsompop, K., & Chaudhari, A. (2023). DDM²: Self-Supervised Diffusion MRI Denoising with Generative Diffusion Models. arXiv preprint arXiv:2302.03018. <https://arxiv.org/abs/2302.03018>
- [44]. Khader, F., Mueller-Franzes, G., Arasteh, S. T., Han, T., Haarburger, C., Schulze-Hagen, M., ... & Truhn, D. (2022). Medical Diffusion: Denoising Diffusion Probabilistic Models for 3D Medical Image Generation. arXiv preprint arXiv:2211.03364. <https://arxiv.org/abs/2211.03364>
- [45]. Dorjsembe, Z., Pao, H. K., Odonchimed, S., & Xiao, F. (2023). Conditional Diffusion Models for Semantic 3D Brain MRI Synthesis. arXiv preprint arXiv:2305.18453. <https://arxiv.org/abs/2305.18453>
- [46]. Wolleb, J., Sandkühler, P., Karimi, D., & Cattin, P. C. (2021). Diffusion Models for Medical Image Analysis. arXiv preprint arXiv:2111.07891. <https://arxiv.org/abs/2111.07891>
- [47]. Todorov, M. I., Müller, S., & Schifferer, M. (2022). Revolutionizing Medical Imaging with Denoising Diffusion Models. *Nature Machine Intelligence*, 4, 843–853.
- [48]. Lyu, Q., Zhang, K., Fu, Y., et al. (2023). 4D Diffusion Models for Dynamic MRI Reconstruction. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14698–14708.
- [49]. Choi, S., Lee, H., & Paik, M. C. (2023). MRI Super-Resolution Using Conditional Diffusion Probabilistic Models. *IEEE Transactions on Medical Imaging*, 42(4), 871–882.
- [50]. Pinaya, W. H. L., Mechelli, A., & Sato, J. R. (2022). Brain Imaging Generation with Latent Diffusion Models. *Medical Image Analysis*, 82, 102605.
- [51]. Zhou, Z., Wang, J., & Bai, J. (2022). Latent Diffusion Models for 3D Medical Image Generation. arXiv preprint arXiv:2211.13867. <https://arxiv.org/abs/2211.13867>
- [52]. Mittal, R., Mewada, R., & Sharma, M. (2023). Denoising Diffusion Models for Multi-Modal Medical Image Synthesis. *Neurocomputing*, 532, 119804.
- [53]. Luo, Y., Chen, M., Zhang, Y., et al. (2022). Multi-Conditional Diffusion Model for Medical Image Synthesis. *IEEE Transactions on Medical Imaging*. <https://arxiv.org/abs/2211.01734>.
- [54]. Ternes, L., Schneider, M., Möller, H. E., & Teschendorff, A. E. (2022). Image Denoising and Reconstruction in MRI Using Denoising Diffusion Models. *Magnetic Resonance Imaging*, 90, 27–39.
- [55]. Haque, A., Bi, J., & Khan, A. (2022). Brain MRI Tumor Segmentation Using Diffusion-Based Image Generation. *Biomedical Signal Processing and Control*, 78, 103920.
- [56]. Tzeng, E., Yang, J., & Sun, W. (2022). Diffusion Models for MR Image Enhancement and Reconstruction. *Computerized Medical Imaging and Graphics*, 97, 102066.
- [57]. Yu, L., Wang, S., Li, C., et al. (2022). A Latent Diffusion Approach for PET-to-MRI Synthesis. *IEEE Journal of Biomedical and Health Informatics*, 26(7), 3314–3323.
- [58]. Schilling, K. G., et al. (2021). Diffusion Models for Tractography in Diffusion MRI: Principles and Challenges. *NeuroImage*, 245, 118702.
- [59]. Han, J., Zhang, C., Huang, Y., et al. (2023). DDPM-MedSeg: Diffusion-Based Medical Image Segmentation. arXiv preprint arXiv:2302.05420. <https://arxiv.org/abs/2302.05420>
- [60]. Bouchacourt, D., Lee, Y., & Alahi, A. (2022). Masked Conditional Diffusion Models for Medical Image Inpainting. *MICCAI Workshops*. <https://arxiv.org/abs/2210.00939>