International Journal of Advanced and Applied Sciences, 12(8) 2025, Pages: 255-272

Contents lists available at Science-Gate

International Journal of Advanced and Applied Sciences

Journal homepage: http://www.science-gate.com/IJAAS.html

Generative AI in neuroscience imaging: A review

Aws I. AbuEid [1], Mohammed Ahmed Elhossiny [2, 3], Marwa Anwar Ibrahim Elghazawy [3], Abdelnasser Saber Mohamed [4], Achraf Ben Miled [4, *], Firas M. Allan [4], Shouki A. Ebad [5], José Escorcia-Gutierrez [6]

[1]*Faculty of Computer Studies, Arab Open University, Kuwait City, Kuwait*
[2]*Faculty of Specific Education, Mansoura University, Mansoura, Egypt*
[3]*Applied College, Northern Border University, Arar, Saudi Arabia*
[4]*Computer Science Department, Science College, Northern Border University, Arar, Saudi Arabia*
[5]*Center for Scientific Research and Entrepreneurship, Northern Border University, Arar, Saudi Arabia*
[6]*Department of Computational Science and Electronics, Universidad de la Costa, Barranquilla, Colombia*

A R T I C L E I N F O

Article history:
Received 22 February 2025
Received in revised form
23 June 2025
Accepted 3 August 2025

Keywords:
Generative AI
Neuroscience imaging
Data augmentation
Anomaly detection
Predictive modeling

A B S T R A C T

Generative AI includes a range of machine learning techniques that model data distributions and generate realistic samples. Methods such as flow-based models, diffusion models, variational autoencoders (VAEs), and generative adversarial networks (GANs) have achieved strong results in various fields. In neuroscience imaging, these techniques can enhance data quality and availability by augmenting datasets, completing missing or noisy data, detecting anomalies, and creating realistic simulations for training predictive models. This review explores the growing role of generative AI in neuroscience imaging, focusing on its applications, benefits, and challenges. It highlights how these models can help overcome data shortages, improve visualization methods, and offer new solutions to persistent problems in the field. By summarizing current research and suggesting directions for future work, this paper aims to support researchers and practitioners in using generative AI to advance neuroscience understanding and improve diagnostic and therapeutic outcomes.

1. Introduction

Neuroscience imaging modalities upgrade the view of the structural details and functioning of the brain, revealing important features of neuronal activity, connectivity, and disorders. Neuroimaging techniques, including magnetic resonance imaging (MRI), functional MRI (fMRI), positron emission tomography (PET), electroencephalography (EEG), and diffusion tensor imaging (DTI), have vastly changed the game for brain researchers and clinicians. MRI and fMRI provide high-quality images of anatomical structures and functional activities (respectively), while PET visualizes metabolic processes (Yen et al., 2023). Functional neuroimaging makes it possible to identify where neural oscillations occur in the brain. Techniques such as EEG provide excellent temporal accuracy, while DTI produces maps of white matter tracts that show brain connectivity. Together, these approaches form a foundation of neuroscience research and are essential for diagnosing and monitoring conditions such as Alzheimer's disease, epilepsy, and psychiatric disorders. However, traditional methods of analyzing neuroscience imaging still have significant limitations. Standard approaches usually depend on pre-defined models, linearity, and hand-crafted features that may not accurately model brain dynamics complexity and individual variability. These methods also have difficulty dealing with the noisy, high-dimensional, incomplete characteristics of imaging data. The need for manual annotation and interpretation, which requires in-depth expertise and a considerable amount of time, limits their scalability at the same time. These challenges underscore the importance of super-resolution computational tools that can reveal hidden patterns, enabling robust, automated analyses and opening up future uses such as deploying generative AI in neuroscience imaging.

Generative artificial intelligence (AI) is a subset of machine learning (ML) methods that attempt to model and understand the underlying data distributions of incoming datasets and generate new examples/pharmaceutical leads that are similar to

* Corresponding Author.
Email Address: ashraf.benmilad@nbu.edu.sa (A. Ben Miled)
https://doi.org/10.21833/ijaas.2025.08.024
Corresponding author's ORCID profile:
https://orcid.org/0000-0003-1256-2900

input data. Unlike traditional discriminative models, which focus predominantly on classification or prediction tasks, generative models aim to generate new points in the data space that resemble the training data (Gupta et al., 2024). These models read the data, understanding all its complex ramifications, like textures, structures, and patterns, which allows them to become a strong solution for image synthesis, data augmentation, and anomaly detection. The main idea of generative AI is to model the data probability distribution and sample from it to create new examples. There have been multiple generative models that have emerged over the years with different methodologies, as shown in Fig. 1. Generative Adversarial Networks (GANs) are among the most popular and transformative approaches. GANs have been widely used in creating realistic images, videos, and even brain imaging data, offering applications in both research and clinical contexts.

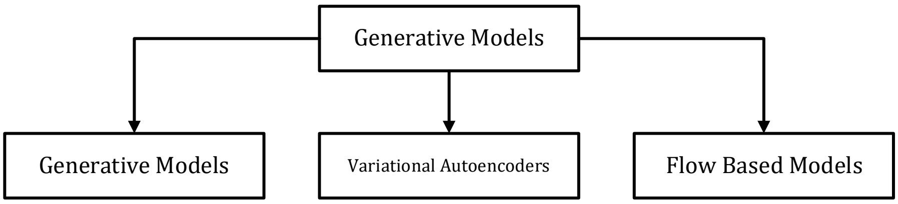

Fig. 1: Generative model types

In contrast, a variational autoencoder (VAE) adopts a probabilistic regime in the family of generative models. It uses an encoder-decoder paradigm where the encoder maps input data into a latent space and the decoder generates data out of that space (Wang et al., 2024a). VAEs tend to make the latent space more consistent by regularizing the learned distribution towards standard distributions (Gaussian, for instance), such that the model transforms input smoothly and generates new smooth outputs. While VAEs excel in interpretability and latent space manipulation, the samples they generate may sometimes lack the fine detail achievable with GANs.

Diffusion models are a newer variety of generative models and have been gaining popularity for their strong performance when it comes to generating high-dimensional data. These models corrupt data with noise in an iterative process and learn to reverse this process (Wang et al., 2024b). Diffusion models are very powerful at generating high-quality images and have been shown to perform well across a variety of tasks, including neuroscience imaging, where their iterative form is naturally suited to model complex and high-dimensional data. Flow-based model provides another generative approach. This enables them to invert the map from sample to data vector with guaranteed likelihood, making flow-based models highly applicable in tasks with explicit probability computations (Jeevan et al., 2024; Song et al., 2024; Fan et al., 2024).

It offers the flexibility of generating data with defined properties and has niche applications in structured datasets. Generative AI has transformed data-driven research by allowing for the generation of high-quality, diverse datasets that can help to overcome issues of data scarcity and improve model interpretability for complex systems. These models have enabled advances in areas that have been limited by small sample sizes and noisy data, such that neuroscience imaging researchers can indeed generate realistic images of the brain, simulate disease progressions, or learn latent representations of neural structures. Generative AI, hence, is a potent aid in neuroscience development.

With traditional approaches in neuroscience imaging facing fundamental limitations, generative models have arisen as a powerful new technology (Hagos et al., 2024). The enhanced dataset can enhance the generalizability of the downstream machine learning models from the point of view of applications like predicting brain segmentation, classification of neural diseases, and forecasting the progress of the disease. Anomaly detection, a critical task in clinical neuroscience, is also an area where generative models prove to be valuable. These approaches essentially work with either handcrafted features or thresholds while defining anomalies. Often, such approaches are inabilities to define subtle patterns and complexity anomalies. These models can learn a normal distribution over brain imaging data, which enables the definition of something being an anomaly as some sort of deviation. For example, autoencoder-based models can reconstruct normal brain images but easily get confused by anomalous patterns that may be useful in detecting tumors, lesions, or degenerative changes. Synthesis of missing data is one of the primary strengths of generative models. Typically, neuroscience imaging datasets consist of missing modalities, motion artifacts, or possibly incomplete scans. These can impute missing data points and reconstruct fully complete datasets with very high fidelity. For instance, GANs and diffusion models have been used to take low-resolution inputs and automatically synthesize high-resolution scans of MRI or generate some missing slices in volumetric brain imaging. This results in the integrity and completeness of incomplete datasets, potentially reducing the need for any repeated scans and thus mitigating patient discomfort.

A generative model can also be applied to generate realistic simulations that may be used for training and testing other models. These simulations may represent detailed neural structures, patterns of disease progression, or patterns of brain activity; therefore, these are high-quality training datasets for large quantities that machine learning algorithms may require. For instance, models can simulate different stages of neurodegenerative diseases, which may help understand the dynamics of the diseases and create predictive models (Fan et al., 2024). Moreover, since the synthetic data produced by such models has no privacy concerns, these can be more freely disseminated and used for collaborative research efforts. Other than addressing these challenges, generative models introduce unprecedented levels of flexibility and innovation in neuroscience imaging. These can learn intricate, nonlinear patterns in the imaging data to discover latent, hidden relationships among the different modalities, and allow for interpretable representations of brain structure and function.

Contributions of the proposed study are as follows: (i) A Comprehensive overview of Generative AI models is presented by highlighting key models (GANs, VAEs, diffusion models, flow-based models) and their relevance to neuroscience imaging. (ii) Applications in neuroscience imaging are discussed and summarized, including anomaly detection to identify irregular brain patterns, reconstruction and synthesis to fill missing data or denoising scans, and to create realistic data for algorithm training. (iii) Analysis of technological trends to discuss advancements in model architecture (e.g., hybrid generative-discriminative approaches) and to explore the tailoring of models for neuroscience-specific tasks, such as disease modeling and brain connectivity mapping. (iv) Challenges in adoption to identify barriers, including data limitations as there is a need for large, high-quality datasets, computational demands for high resource requirements, and interpretability issues to hinder trust and clinical integration. (v) Impact on Neuroscience research and clinical practice to emphasize the transformative potential for modeling disease progression, advancing personalized diagnostics, and improving the efficiency of research workflows.

By addressing these aspects, the review contributes a well-rounded perspective on how generative AI is reshaping neuroscience imaging, while also laying the groundwork for future advancements.

The rest of the paper is summarized as follows: Section 2 focuses on different types of generative models and their application in neuroscience imaging, followed by Section 3, which focuses on applications in neuroscience; Section 4 discusses challenges and future directions, followed by a conclusion in Section 5.

2. Generative models in neuroimaging

The central focus of this section is on the various types of generative models and their specific applications in neuroscience imaging.

2.1. Generative adversarial networks (GANs)

Generative Adversarial Networks (GANs), introduced by Goodfellow et al. (2014), have become a prominent approach for generative modeling due to their ability to produce high-quality synthetic data. The objective of GANs is formulated as a minimax game with the following loss function:

$$\min_{G}\max_{D} E_x \sim p_{data}(x)[\log D(x)] + Ez \sim p_z\,(z)\left[\log\left(1 - D\big(G(z)\big)\right)\right] \quad (1)$$

where, x is a real data sample, z is a random noise vector sampled from a prior distribution $p_z\,(z)$ and $G(z)$ is the synthetic data generated by G.

GANs have been applied in various fields, including image synthesis, image-to-image translation, super-resolution, and data augmentation. Despite their success, GANs often suffer from instability during training and mode collapse, where the generator produces a limited diversity of samples. Numerous variants, such as Wasserstein GANs (WGANs) and StyleGANs, have been proposed to address these limitations. A GAN consists of two neural networks: A generator G and a discriminator D, which are trained simultaneously through adversarial training. The generator aims to produce data samples that resemble real data, while the discriminator attempts to distinguish between real and generated samples, as shown in Fig. 2. This adversarial framework enables GANs to generate highly realistic data across various domains, including neuroimaging.

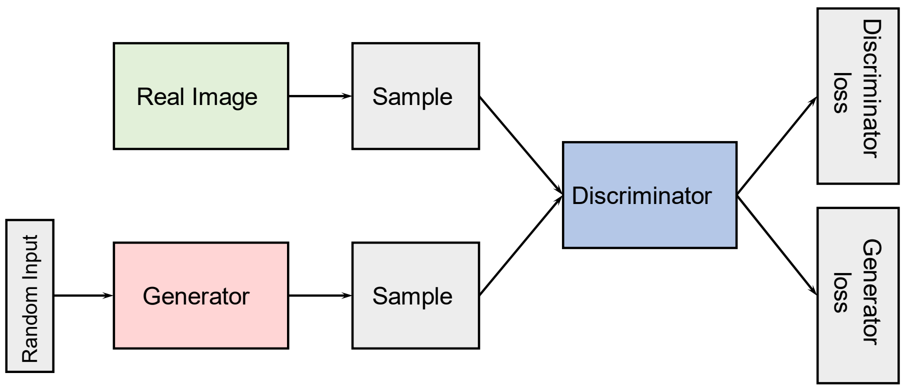

Fig. 2: Structure of GAN (Yilmaz and Korn, 2024)

The generator starts with random noise and learns to map it to the data distribution of the target dataset. It achieves this through a series of transformations, progressively refining the synthetic output to resemble real samples. The discriminator, on the other hand, acts as a binary classifier that distinguishes between real and synthetic data. GAN training relies on a loss function derived from game theory, commonly the minimax objective (Ledig et al., 2017). While the discriminator seeks to maximize its classification accuracy, the generator seeks to reduce the likelihood that it would accurately identify its outputs as fraudulent.

Training GANs is inherently difficult as these easily cause instability, mode collapse, in which the generator generates fewer variabilities and vanishing gradients. Techniques like Wasserstein GANs, feature matching, and progressive growing have thus been proposed that make the GAN much more stable. Applications for GANs have indeed been made in neuroimaging, such as realistic images of brain generation, and synthesis of imaging modalities not available in one's system (Wang et al., 2023a; Sabuhi et al., 2021). Generative models have led to transformative applications in neuroimaging by overcoming longstanding challenges in data quality, completeness, and analysis. Table 1 lists the overview of GANs applications in neuroimaging, summarizing relevant literature with results, techniques, remarks, strengths, and weaknesses.

2.2. Variational autoencoders (VAEs)

Variational Autoencoders are generative models of data. It learns to map data to a latent space, which allows generating new samples that resemble the training data (Molnár and Tamás, 2024). VAEs are much better than basic autoencoders in that they apply probabilistic inference. The latent space in this case follows a set distribution, which is multivariate Gaussian as shown in Fig. 3. This would make VAEs effective in generating smooth, continuous variations in data, suitable for image synthesis and learning representation in neuroimaging.

The two elementary building components of VAE architecture include the encoder and decoder. An encoder transforms the input into latent form in estimating posterior $q(z/x)$. z defines a latent variable, while x symbolizes the input. In doing this, it provides the mean as well as the variance for the Gaussian in form describing the z (Chadebec et al., 2022). From this distribution, the decoder takes a sample z, decodes it back into the data space, and rediscovers the input. It will train by estimating in that process to $p(x/z)$. The recreation of its output shall make it look like the real data. Improving a loss function VAE's means kind that has two components: Its reconstruction loss plus its KL divergence. Reconstruction loss allows the decoder to generate outputs that closely match the original data. Typically, it is computed by metrics like mean squared error. The KL divergence maintains the structure in latent space because it ensures $q(z|x)$ is equal to some prior distribution $p(z)$. It therefore makes latent representations smooth and meaningful. The total loss is given in Eq. 2:

$$L = E_{\{q(z|x)\}[\log p(x|z)]} - KL\big(q(z|x)/p(z)\big) \tag{2}$$

VAEs are strong and easy to understand, making them liked for tasks such as learning shorter forms of neuroimaging data, modeling how diseases get worse, and creating realistic brain scans for research and clinical use.

Variational autoencoders have been recognized as one of the key technologies in neuroimaging to overcome the problems of high-dimensional data, small dataset sizes, and the requirement of interpretable representations. Their ability to learn probabilistic mappings and compact latent spaces has opened innovative applications in the field. One of the most important applications of VAEs is dimensionality reduction. Neuroimaging data, such as MRI, fMRI, and PET scans, are often high-dimensional and computationally expensive to process (Wei and Mahmood, 2020). This latent representation can then be used further for classification, clustering, or regression; it simplifies complex patterns of the brain. For example, VAEs have been employed in the reduction of dimensionality to study networks of connectivity in the brain; it does indeed show some interesting patterns compared with control conditions. VAEs are used widely in neuroimaging for data augmentation purposes; datasets are always small due to the extremely high cost of collecting and the rather low accessibility to them. This synthetic data will improve the diversity of training datasets of machine learning models and enhance the robustness and generalizability of such models. For example, VAEs have been used in synthesizing brain MRIs with different anatomical structures or pathological features for tasks in tumor detection and segmentation.

Another significant application of VAEs is in latent space analysis, helping to understand the complex data of brain images. Latent space, learned by VAEs, demonstrates a much clearer and simpler representation of the input data that shows its most important features. Researchers may change this space to observe the relationship between different parts of the brain, or conditions and stages of the disease. For instance, VAEs were used in developing brains with the disease getting worse over time to point out how the disease is moving at the connecting points within the hidden space, hence painting a clearer picture of changes over time in the brain's structure. Such insights may help in the tailoring of treatment or pinpoint what causes disorders of the brain.

VAEs also help combine different types of brain scans, like structural MRI with functional MRI or PET scans. By aligning the hidden features of these scans, VAEs make it easier to analyze them together, revealing patterns that are hard to find with just one type of scan.

Table 1: Generative adversarial networks (GANs) in neuroimaging

Reference	Technique	Application	Results	Remarks	Strengths	Weaknesses
Bowles et al. (2018)	3D GAN for brain tumor segmentation	Tumor segmentation	Improved tumor segmentation accuracy	Introduced a 3D GAN for better spatial understanding	Handles 3D spatial data effectively	High computational cost
Han et al. (2019)	CycleGAN for cross-domain image translation	CT to MRI image synthesis	High-quality MRI synthesis with structural similarity indices	Useful for reducing dependency on multi-modal scans	Eliminates the need for both CT and MRI during training	Limited generalization across datasets
Nimeshika and Subitha (2024)	Conditional GANs generate realistic synthetic samples for minority classes	Utilizes split federated learning (SFL) to enable collaborative training without sharing sensitive medical data.	Accuracy = 83.54%	Integration of SFL and cGANs to address challenges in medical classification for decentralized, imbalanced datasets.	Helps in understanding disease progression	Increase training time and computational resources
Shin et al. (2018)	Super-resolution GAN for upscaling low-resolution MRI scans	MRI enhancement	Enhanced MRI quality, enabling better diagnosis	Provides high-resolution MRI images from low-quality scans	Useful in low-resource settings	Risk of generating artifacts
Nie et al. (2017)	3D GAN for missing data imputation	Reconstruction of incomplete MRIs	Generated plausible reconstructions of missing regions	Demonstrates GAN's ability to reconstruct incomplete brain images	Effective for recovering incomplete neuroimaging data	Challenging to train and optimize
Chen et al. (2023)	Dual Multilevel Constrained Attention GAN (DMCA-GAN)	Hippocampus segmentation	Achieved a Dice coefficient = 90.53% on the MSD dataset, outperforming the baseline by 3.78%.	Significant improvement in segmentation accuracy	Balances noise suppression and feature enhancement	Computationally intensive
Xu et al. (2019)	Semi-Supervised Attention-Guided CycleGAN (SSA-CycleGAN)	Data augmentation in MRI images	Generated realistic synthetic tumor/normal images by adding/removing tumor lesions.	SSA-CycleGAN for improving medical image classification tasks	Effectively enhances the model's focus on important image details through attention-guided modules.	High computational resources
Sajjad et al. (2021)	Deep convolutional generative adversarial network (DCGAN)	Data augmentation	Accuracy = 72%	Synthesis quality and classification performance are notable	Successfully synthesized realistic PET images for all three stages of Alzheimer's disease	Accuracy = 72%
Tan et al. (2024)	Deep Convolutional GAN (DCGAN)	Augment fMRI functional connectivity (FC) data for classifying altered brain networks.	Significant improvement in classification accuracy for major depressive disorder (MDD) identification	Generates realistic synthetic FCs with structural patterns resembling real data	Limited exploration of potential biases introduced by synthetic data	Data generalizability in neuroscience imaging, for tasks with limited datasets like MDD classification
Hwang and Shin (2024)	Conditional diffusion model for image-to-image translation	Generated realistic brain images capturing Alzheimer's progression	Validated the superiority of multi-modal datasets	High-quality image generation and progression modeling	Promising technique with the potential to advance Alzheimer's diagnosis	Extensive computational resources

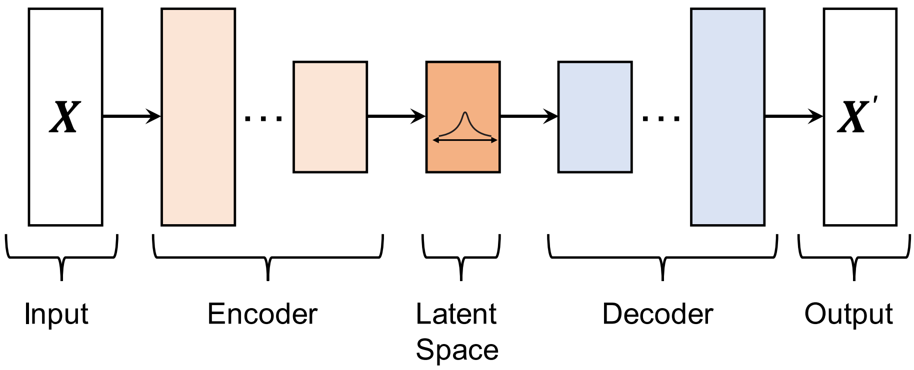

Fig. 3: Variational autoencoder (Singh and Ogunfunmi, 2021)

This is especially useful for understanding complicated things like how different parts of the brain connect or how structural and functional problems work together. VAEs are useful and effective but also come with challenges in neuroimaging. The data these create sometimes misses the small details that other models like GANs can capture. Also, it is difficult to adjust the latent space to make it easier to understand. Still, their ability to reduce data size, improve datasets, and give meaningful latent representations makes VAEs very important for progress in neuroimaging research and clinical practice. VAEs have been highly successful in neuroimaging, providing strong methods for dimensionality reduction and exploration of the latent space. Their probabilistic formulation allows for the generation of realistic and diverse samples of data that are most useful in application domains in which the lack of abundant data with sufficient variability becomes a problem (Yin et al., 2025). For example, VAEs have been applied very effectively to synthesize neuroimaging data such as MRI scans to improve the performance of machine learning models. Also, their skill in capturing important hidden representations has helped in studying complicated events like disease development and combining different types of data, giving useful information about brain structure and function.

However, VAEs have very important problems. The most significant problem is blurry images. Indeed, the quality of images synthesized by VAEs can be low since they base their representation on the Gaussian distribution within the latent space and a likelihood-based objective function. This naturally leads to blurred details of images, which are crucial in neuroimaging for proper assessment of small anatomical or pathological details. Approaches toward solving this problem include using perceptual losses or combining VAEs with GANs. However, it's challenging to get high-quality, detailed reconstructions. In addition, VAEs are sensitive to hyperparameters. Their balance between reconstruction loss and the KL divergence term is very critical. If not done right, these can cause bad results like over-regularized latent spaces or very simple representations. Furthermore, while the interpretability of VAEs is achieved through latent space analysis, the interpretation of data sets may be highly complex, requiring additional techniques for meaningful insights. Yet the strengths of VAEs, including flexibility, probabilistic framework, and generalization capability, make them a promising tool in neuroimaging applications, and ongoing research seeks to address the shortcomings. The literature in the field of neuroimaging has widely explored and applied the variational autoencoder due to its capability across different applications.

In a study by Bit et al. (2024), VAEs were used to analyze neuroimaging data from patients with schizophrenia and Alzheimer's disease. By reducing the dimensionality of structural MRI scans, the VAE extracted compact latent representations, which were then analyzed to identify patterns associated with these disorders. The study demonstrated that VAEs could reveal disease-specific features, aiding in diagnostic and prognostic assessments. A study by Nalepa et al. (2019) employed VAEs to generate synthetic brain MRI scans for training deep learning models. The augmented datasets improved the performance of models in tasks such as tumor segmentation and brain region classification. This highlighted the role of VAEs in overcoming the limitations of small sample sizes in neuroimaging studies. Zhou et al. (2019) used VAEs to model the progression of Alzheimer's disease. By interpolating between points in the latent space, the VAE provided a smooth representation of disease evolution, visualizing the gradual transition from normal brain structures to pathological states. This application demonstrated the potential of VAEs in understanding the temporal dynamics of neurodegeneration. In a multi-modal application, Reaungamornrat et al. (2022) utilized VAEs to synthesize missing imaging modalities, such as generating functional MRI data from structural MRI scans. The study highlighted how VAEs could integrate and reconstruct complementary information across different neuroimaging modalities, enhancing analyses in resource-constrained settings. VAEs have also been employed to harmonize data from different imaging centers. For example, Abbasi et al. (2024) applied VAEs to correct site-specific variations in structural MRI data collected from multiple institutions. This enabled the researcher to combine data across sites, improving the generalizability of their findings. A study by Chatterjee et al. (2022) used VAEs to detect anomalies in brain MRIs. By learning a latent space that represented normal brain structures, the VAE flagged abnormalities such as lesions or tumors as regions that deviated significantly from the learned normal distribution. These studies illustrate the versatility of VAEs in tackling challenges like data scarcity, high dimensionality, and modality-specific limitations in neuroimaging. While those have clear limitations, such as blurry reconstructions, their successes across these applications underscore their value in advancing neuroimaging research and clinical practices. Table 2 summarizes the key insights on the use of Variational Autoencoders (VAEs) in neuroimaging based on current literature, including results, techniques, strengths, weaknesses, and remarks.

Table 2: Variational Autoencoders (VAEs) in neuroimaging

Reference	Techniques	Results	Strengths	Weaknesses	Remarks
Rais et al. (2024)	VAEs for data augmentation, segmentation, and classification	Improved dataset balance and enhanced segmentation/classification accuracy	Effective for small, imbalanced datasets and realistic synthetic data generation	Limited diversity in generated samples	Comparison with GANs shows VAEs generate realistic data but need better sample diversity
Wei and Mahmood (2020)	Representation learning with VAEs in medical imaging	Improved tumor segmentation and structural analysis in MRI images	Strong feature extraction capabilities for high-dimensional data	Sensitive to hyperparameter tuning; prone to mode collapse	Effective for unsupervised learning tasks in medical imaging
Sidulova, and Park (2023)	Conditional VAEs for functional connectivity modeling	Enhanced prediction of neurological biomarkers from fMRI data	Robust representation of complex spatial-temporal patterns	Computationally expensive, limited interpretability	Useful in identifying disease progression patterns
Qiang et al. (2020)	Deep Variational Autoencoder (DVAE)	Achieved state-of-the-art classification accuracies	Learned representations were interpretable and meaningful, with functional brain network (FBN) patterns organized hierarchically	Dataset dependency with computational complexity	Highlights the potential of DVAE for addressing key challenges in neuroscience imaging

2.3. Diffusion models

Diffusion models have emerged as a powerful class of generative models, particularly in the field of neuroimaging. These models operate on the principle of gradually adding and then removing noise from data, enabling the generation of high-quality images (Cao et al., 2024). Mathematically, this is often modeled as a Markov chain, where each step depends only on the previous one. The reverse diffusion process is the key to generating new data. This step is followed by denoising training, wherein an image neural network steps its way into de-noising data effectively undoing the diffusion process. Diffusion processes, hence, allow iteration such that their models keep re-refining outputs with great production of highly detailed and realistic images. These models find wide applicability in neuroimaging studies. Their enhanced images resolve noise, hence increasing the possibility of proper identification of various neurological disorders. These models can be further used to create synthetic datasets that aid in the training of machine learning models. This is particularly helpful in neuroimaging, where quality data may be limited. Furthermore, diffusion models can assist in the diagnosis of diseases and tracking disease progression by improving the clarity and accuracy of imaging data. It will help doctors generate images that are detailed and realistic so that they can better understand the progression of neurological diseases. Diffusion models utilize forward and reverse diffusion process principles to create high-quality images in neuroimaging.

Their potential to enhance image quality, synthetic data creation, and better diagnosis of diseases makes them very significant in the field of neuroscience. Diffusion models have significantly advanced the field of neuroimaging by providing robust applications such as high-quality image generation, as demonstrated in Fig. 4.

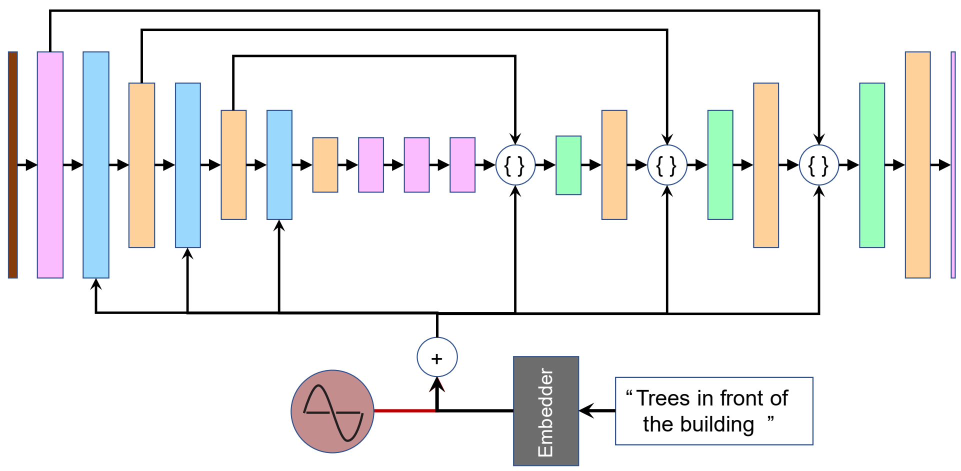

Fig. 4: Diffusion model

One of the main applications is the generation of high-quality images. Neuroimaging often captures intricate details of brain structures, which are essential for accurate diagnosis and research. Diffusion models denoise and enhance the resolution of these images to create clear and sharper images. This can be particularly useful in the clinic where good-quality images are important to spot small abnormalities and make a treatment decision. Acquiring very large and diverse datasets can be very challenging in neuroimaging because of ethical, financial, and logistical reasons. Diffusion models thus augment the size of the dataset, helping researchers develop more accurate and reliable algorithms for different neuroimaging tasks, like disease classification and segmentation. In addition, diffusion models help in reconstructing damaged or partially missing images without losing valuable information due to technical artifacts or limits. Reconstruction of images, therefore, improves the quality of neuroimaging data as a whole, thus becoming more reliable for use in clinical applications and research. Generally, the diffusion models essentially contribute toward advanced neuroimaging as these generate high-quality images, augment datasets, and reconstruct incomplete data, thereby providing better diagnostic accuracy and research outcomes. A tremendous amount of success has been recorded in the field of neuroimaging from the diffusion model, as it improves image quality, along with enhancing data. One of the greatest strengths of diffusion models is the creation of noise from real-noise observations to yield a more realistic output image of greater resolution. This feature ensures effective enlargement or enhancement of the image clarity needed for making proper diagnoses or conducting neuroscientific research. Their ability to generate synthetic datasets for training machine learning algorithms makes them even more useful. Nonetheless, the same diffusion models are not without drawbacks, even if the researchers who have used them have noted the following successes. One of the major concerns is the computational complexity of these models. During the training of diffusion models, a lot of data is involved, and many algorithms need to be executed, which costs a lot of computational power and time. Moreover, error rates do call for refinement and testing on many parameters that diffusion models claim require elaborate and extensive parameters, which can take time and form a technical challenge. The second is that if the input data is of bad quality, or it has some pre-existing biases, then a diffusion model may not output proper values that are not infected by these biases. However, the details may be overlooked when using diffusion models, while image reconstruction and denoising may be highly accurate, depending on the medical application. Although useful to enhance the quality of the image and also augment the data set for neuroimaging applications, diffusion models also incur heavy computational costs along with dependency on good-quality input data. There may exist scope to overcome the constraints and eventually make these even more fruitful and practical to apply to neuroscience and brain function applications. Yen et al. (2023) discussed recent developments in noninvasive functional neuroimaging methods, including fMRI and EEG. It highlights the role of advanced neuroimaging techniques, such as diffusion tensor imaging (DTI) and transcranial electrical stimulation (TES), in studying brain connectivity and potential treatments for neurological disorders. The chapter provided an overview of quantitative computational methods for analyzing neuroimaging data, including diffusion MRI data. It discusses how methods developed for traditional scalar structural neuroimaging data have been extended to diffusion MRI data, allowing the study of the brain's connection structure (O'Donnell and Schultz, 2015). These examples illustrate the diverse applications and advancements of diffusion models in neuroimaging, showcasing their potential to improve diagnostic accuracy and research outcomes.

Application of diffusion models in neuroimaging is presented in Table 3 with details on techniques, results, strengths, and weaknesses. Diffusion models excel in reconstructing high-resolution and temporally consistent neuroimaging data. These models offer flexibility for tasks like super-resolution and imputation by leveraging learned noise distributions. These are effective in settings with limited high-quality data, such as ultra-high-field MRI applications. Whereas training diffusion models demands significant resources.

Table 3: Application of diffusion models

Reference	Techniques	Results	Strengths	Weaknesses
Yuan et al. (2024)	Conditional denoising diffusion probabilistic models for longitudinal MRI imputation	Enhanced recovery of missing MRI scans from adjacent timepoints; SSIM and NMSE improved compared to GAN-based models	Robust handling of missing data; effective use of temporal information	Requires computationally intensive training and fine-tuning
Yoon et al. (2024)	Diffusion model-based generative AI (d3T) and CNN-based model (c3T) for superresolution (SR) of 1.5T MR images to emulate 3T images	AD classification using 3T outperformed 1.5T in accuracy, sensitivity, and specificity	Significantly enhanced image quality and volumetric accuracy. Improved diagnostic and prognostic performance for AD and MCI	Dependence on advanced computational resources for model training and inference. Potential for reduced generalizability to diverse datasets outside ADNI1
Gajjar et al. (2024)	Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), Diffusion Models, and DenseNets for classification	Diffusion models generated non-dementia images with FID = 92.46; GANs excelled in dementia images (FID = 178.53)	Diffusion models excel in generating non-dementia images with low FID	Potential dataset biases not addressed

Noise modeling and parameter tuning can be challenging, especially in clinical contexts, and understanding and validating the outputs of diffusion-based models in medical settings remains an active research area. Diffusion models are becoming indispensable in neuroimaging, particularly for enhancing image resolution, addressing missing data, and modeling disease progression. Future directions may include integrating diffusion models with real-time clinical pipelines for diagnostic and therapeutic applications.

2.4. Flow-based models

Flow-based models, unlike GANs, are explicit generative models that provide exact likelihood estimation shown in Fig. 5. Table 4 lists the summary of key research studies in this domain, highlighting different techniques, results, and remarks. These models achieve this by modeling the data distribution using a series of invertible transformations.

The core idea behind flow-based models is to learn a bijective mapping f between the input data x and a latent variable z drawn from a simple prior distribution (e.g., a standard normal distribution). The probability density function of the input data x can be computed using the change of variables formula:

$$p(x) = p(z)\left|det\left(\frac{\partial f}{\partial x}\right)\right|^{-1} \quad (3)$$

By designing f to be invertible and its Jacobian determinant to be efficiently computable, flow-based models enable exact likelihood computation and sampling. Prominent examples of flow-based models include RealNVP, Glow, and NICE. Flow-based models are a category of explicit generative models that enable precise likelihood estimation by transforming data distributions through a sequence of invertible functions. These models establish a bijective mapping between input data and a latent variable drawn from a predefined distribution, such as a standard normal distribution. The probability density of the input data is determined using the change of variables formula. A key feature of flow-based models is their ability to ensure invertibility and efficient computation of the Jacobian determinant, which facilitates exact likelihood estimation. Prominent examples include RealNVP, Glow, and NICE.

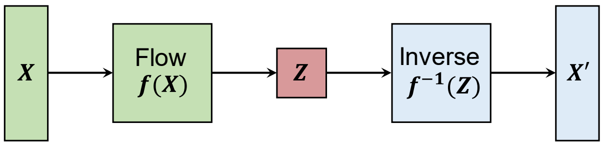

Fig. 5: Flow-based model (Wang et al., 2023b)

Unlike other generative models such as GANs and VAEs (Tomczak, 2020; Gupta et al., 2025), these models maintain full data invertibility, allowing transformations to be reversed precisely to recover the original input. This property enables accurate probability estimation, setting them apart from alternative approaches. Invertible transformations play a central role in flow-based models. While other generative models apply transformations that are difficult to reverse, flow-based architectures ensure each transformation remains computationally feasible to invert. Some commonly used invertible transformations include coupling layers, autoregressive transformations, and 1x1 invertible convolutions. Additionally, these models' ability to compute exact likelihood values makes them valuable for applications requiring precise probability evaluations.

The flexibility of flow-based models allows them to capture intricate dependencies in data by stacking multiple layers of invertible transformations. Each layer contributes to enhancing the model's expressiveness, making it well-suited for tasks such as data augmentation and synthetic data generation. By offering accurate likelihood estimates, these models can reveal data patterns that may indicate neurological conditions, making them particularly useful in neuroimaging applications (Choi and Sunwoo, 2022; Bacon et al., 2024). In neuroimaging, obtaining diverse datasets is often constrained by ethical, logistical, and financial challenges. Synthetic data generated through flow-based models can enhance machine learning models by increasing the volume and diversity of training examples, ultimately improving model robustness. These models assist in various neuroimaging tasks, including disease classification, segmentation, and density estimation. By estimating probability densities of neuroimaging data, flow-based approaches help detect anomalies and outliers, aiding in medical diagnoses. Another significant application of flow-based models in neuroimaging is image denoising and enhancement. Due to their ability to compute precise likelihoods, these models effectively reduce noise and enhance image resolution, resulting in clearer and more detailed scans.

High-quality neuroimaging data enables healthcare professionals to detect minute abnormalities, leading to more accurate diagnoses and research outcomes. Despite their advantages, flow-based models have certain limitations. Training these models requires substantial computational resources and expertise in advanced machine learning techniques, making them less accessible to researchers unfamiliar with deep learning. Additionally, poor-quality training data can introduce biases, affecting model reliability.

Table 4: Flow-based techniques

Reference	Techniques	Results	Strengths	Weaknesses
Ahmadi et al. (2024)	Euclidean Ricci Flow, Covariance Matrices	Introduced landmark-free methods for Alzheimer's diagnosis, using multi-modal surface properties for cortical morphometry	High accuracy in cortical surface analysis, non-reliant on manual landmarks	Requires sophisticated computational methods, not yet widely implemented
Gong et al. (2023)	Generative AI, Convolutional Neural Networks, Cross-Modality Imaging	Generative models improved brain image analysis tasks like segmentation, classification, and super-resolution	Comprehensive approach to multi-modal imaging, enhancing resolution and classification	Still in the early research stages, challenges in integrating with clinical practices
Sun et al. (2019)	Flow-based generative models (DUAL-GLOW) for PET image synthesis from MRI	Successfully modeled brain FDG-PET hypometabolism changes as a function of age	Novel formulation of conditional distribution between MRI and PET latent codes	Significant computational complexity
Zhen et al. (2021)	Three invertible layers for manifold-valued data in flow-based generative models. Two-stream GLOW model for modality transfer/translation between manifold-valued measurements (e.g., orientation distribution functions (ODF) and diffusion tensor images (DTI))	Successful modality transfers: Accurately reconstructed brain images of ODF from DTI, demonstrating reliable performance despite lower angular resolution in DTI. Fast acquisition: DTI (5× faster acquisition time than ODF) with acceptable trade-offs in angular resolution	Expands generative models to non-Euclidean/manifold-valued data (particularly in brain imaging). High accuracy and efficiency for modality transfer. Theoretical contribution with invertible layers for manifold data. Scalable approach with potential broader applications	Sparse literature and limited research on generative models for manifold-valued data make it a new and niche area. The approach may require further optimization for larger datasets or more complex modalities beyond brain imaging. Potential limitations in generalizability to other fields outside brain imaging
Bui et al. (2020)	Invertible architecture for unpaired image-to-image translation using temporal information and deformation fields to guide translation in medical images.	Competitive performance in MSE, PSNR, and SSIM. Achieved on datasets HCP, MRBrainS13, and Brats2019. Synthesized images are realistic and consistent	No additional loss functions are needed due to the invertible flow-based architecture. Utilizes temporal constraints for improved image consistency. High performance on standard datasets	Lack of guarantee for a unique one-to-one mapping between image domains. It may be limited by the complexity of handling large-scale volumetric data
Wei et al. (2023)	SOFNet, utilizing an optical flow-based encoder–decoder backbone for MRI data augmentation	Achieved significant enhancement in super-resolution quality. Surpassed other SISR methods in feature completion and clarity of interpolated slices	Effectively addresses larger slice gaps (4.2 mm to 6.0 mm)	Potential computational complexity of optical flow and encoder-decoder models for larger datasets

Scaling these models to handle extremely large or high-dimensional datasets also presents challenges, potentially limiting their usability in certain neuroimaging applications. Flow-based models provide substantial benefits in neuroimaging, including precise likelihood estimation, data augmentation, and image refinement. However, their computational demands, technical complexity, and sensitivity to data quality must be addressed for broader adoption. Studies have explored their role in reconstructing neuroimaging data and mitigating motion artifacts, highlighting their potential to improve diagnostic accuracy. Systematic reviews have further assessed how AI-driven neuroimaging analysis, including flow-based methods, enhances disease classification and lesion segmentation. Other generative approaches with promising applications in neuroimaging include autoregressive models, transformer-based architectures, and conditional generative models. Each of these methods offers distinct advantages for enhancing neuroimaging data, contributing to improved diagnostic and research capabilities. Flow-based models have also been employed in advanced mathematical techniques such as Ricci flow for brain surface analysis, particularly for identifying conditions like Alzheimer's disease. These models optimize surface features for classification tasks without relying on manual landmark identification. Generative models have significantly advanced brain image computing by improving classification and cross-modality imaging. While these models offer substantial benefits, challenges remain in their clinical integration and computational efficiency. As research progresses, refining these techniques will enhance their applicability in neuroimaging, further supporting diagnostic and analytical advancements in the field.

2.5. Comparative analysis and relevance

Flow-based models have several advantages over GANs, particularly in their ability to estimate likelihood directly. This feature is especially useful in applications that require accurate probability assessments. Additionally, their invertible nature makes them well-suited for tasks like image generation, anomaly detection, and density estimation. In contrast to GANs, which frequently encounter challenges such as mode collapse and unstable training, flow-based models exhibit more consistent training behavior. However, this stability comes with a trade-off—these models typically require a greater number of parameters and significantly more computational power. The balance between model complexity and generative performance is a key factor in determining their practical usability.

For a well-rounded discussion, it is important to present flow-based models not only in terms of their strengths but also by addressing their limitations compared to GANs. A structured comparison of their benefits and drawbacks will improve clarity and offer a deeper understanding of their role in generative modeling.

This study conducts a comparative analysis of GANs, Diffusion Models, and Flow-Based Models using benchmark datasets such as MNIST and CIFAR-

10. The evaluation focuses on essential image reconstruction metrics, including Peak Signal-to-Noise Ratio (PSNR) for image quality assessment, Structural Similarity Index (SSIM) for measuring perceptual similarity, and Fréchet Inception Distance (FID) for comparing the distributions of real and generated images. Findings reveal that while GANs generate high-quality images, they often suffer from instability and mode collapse. Diffusion models provide better training stability and produce high-fidelity images, but demand significant computational resources. Flow-based models, despite their ability to generate precise outputs, are computationally expensive and complex to implement.

Table 5 provides a comprehensive overview of the experimental findings, comparing the performance of various models on the MNIST and CIFAR-10 datasets using multiple evaluation criteria. The results illustrate the balance between image quality, training consistency, and computational efficiency. This analysis helps researchers understand the advantages and drawbacks of each generative model, assisting them in choosing the best-suited method for their particular needs.

Table 5: Comparative analysis of GANs, diffusion models, and flow-based models

Model	PSNR (dB)	SSIM	FID	Training stability	Computational complexity
GAN (MNIST)	27.3	0.89	12.5	Moderate	Moderate
Diffusion (MNIST)	30.1	0.92	10.2	High	High
Flow-Based (MNIST)	28.5	0.90	11.0	High	High
GAN (CIFAR-10)	24.6	0.82	19.6	Moderate	Moderate
Diffusion (CIFAR-10)	26.9	0.87	17.2	High	High
Flow-Based (CIFAR-10)	25.5	0.85	18.0	High	High

3. Applications of generative models in neuroscience

Generative models have brought significant advancements to neuroimaging by tackling key challenges such as small datasets, image quality improvement, disease analysis, and personalized treatment. Table 6 outlines the various applications of these models in neuroimaging, highlighting the techniques used and their advantages. By generating synthetic data, these models help expand training datasets, enhancing the accuracy and reliability of machine learning algorithms.

Approaches like GANs and Variational Autoencoders (VAEs) have contributed to higher image resolution, noise reduction, and artifact removal, leading to more precise diagnostics. Furthermore, generative models support disease simulation, tailored treatment strategies, anomaly detection, and the customization of brain atlases, greatly benefiting both clinical practice and neuroscience research.

Table 6: Applications of generative models in neuroimaging

Application	Description	Key techniques	Key benefits
Data augmentation	Generative models generate synthetic neuroimaging data to expand training datasets, improving machine learning model accuracy and generalization	Generative Adversarial Networks (GANs), VAEs	Helps overcome limited datasets, improves model robustness and accuracy, supports diverse and generalizable algorithms
Image enhancement and restoration	Enhances neuroimaging data through techniques like denoising, super-resolution, and artifact removal	GANs, Convolutional Autoencoders (CAEs)	Improves image quality, enables clearer diagnosis by removing noise and artifacts, and enhances low-resolution images for better anatomical visualization
Disease modeling and simulation	Simulates brain disease progression and treatments for personalized predictions and therapeutic strategies	Deep Learning, Digital Twins, Variational Autoencoders	Allows for realistic disease simulations, predicts outcomes, helps in testing treatment strategies, and develops new therapies
Brain atlas creation and individualization	Creation of personalized brain atlases based on individual neuroimaging data to reflect unique brain structures and functions	VAEs, GANs, Flow-based Models	Enhances accuracy and resolution in brain mapping, and allows personalized insights for research and clinical use, particularly in neurosurgery and diagnosis
Anomaly detection	Identifies abnormal patterns in neuroimaging data, such as structural or functional deviations, indicating possible diseases like Alzheimer's, tumors, etc.	GANs, VAEs, Flow-based Models	Detects subtle anomalies that may not be visible to the human eye, improves diagnostic accuracy, and supports early disease detection
Personalized medicine	Provides tailored diagnostics and treatment plans based on individual neuroimaging data, facilitating personalized care and monitoring disease progression	Digital Twins, Personalized Simulations, GANs	Enables precise diagnosis, predicts treatment outcomes, facilitates tailored therapies, and improves patient outcomes by simulating interventions

3.1. Data augmentation

Generative models greatly ease the burden of limited neuroimaging datasets by creating high-quality synthetic data that mimics real images. The augmenting data expands the training set and, therefore, increases the robustness and accuracy of the machine-learning models.

Through the provision of diverse and large training examples, generative models enable the development of more reliable and generalizable algorithms, thus furthering research and clinical applications in neuroimaging.

3.2. Image enhancement and restoration

Generative models play a crucial role in enhancing and restoring neuroimaging data through applications in denoising, super-resolution, and artifact removal.

Denoising: In most cases, the neuroimaging data is noisy because of the weak signal, subject motion, or limitations in the imaging equipment. These models can effectively learn the difference between a signal and noise, and maintain all the relevant anatomical features while eliminating any undesirable noise. This enhancement is very important for reaching a correct diagnosis and further analysis of the results. High resolution is critical in Neuroimaging because even small structural and functional alterations to the brain can be observed. One benefit of a generative model is that it can improve on low-quality images by creating high-quality images. The technique, called super-resolution, uses computer algorithms to teach models how to accurately enhance images at a microscopic level while preserving the internal structure. In the case of super-resolution, GAN has been more useful in generating higher-resolution scans from low-resolution scans, especially for visualizing the structures of the brain. As it has been stated above, neuroimaging data can be contaminated by different types of artifacts, including motion artifacts, scanner distortions, and numerous other disturbances. For these artifacts, generative models are capable of removing them, which in turn leads to more accurate images. For example, a study applied the GANs to fix motion artifacts in MRI scans based on the fact that GANs can learn to create images that look like real images without artifacts. This capability enhances the standard of the data, making it more appropriate for use in the clinical and research domains. Thus, generative models significantly enhance neuroimaging data through effective denoising, super-resolution, and artifact removal techniques. These applications improve the quality and reliability of neuroimaging data, facilitating better diagnosis, research, and understanding of neurological conditions.

3.3. Disease modeling and simulation

Generative models are already rapidly transforming disease modeling and simulation in neuroscience by providing exceptionally realistic renditions of diseases inside the brain. Development of these diseases into reality is mimicked, whereby knowledge of the mechanisms occurs, and potentially better treatments may be found. For instance, deep generative models can synthesize the personal digital twins of patients given data on the particular person, thereby allowing for the personalized prognosis of the disease's trajectory (Seiler and Ritter, 2025). This way might help scientists as they study and develop many approaches to treatment and estimate prognosis with personalized and individualized medical treatments. Furthermore, the generative models enable the qualification of the impact of any given intervention on the structure and function of the brain. This capability will be more useful, especially when researchers are testing new drugs and or treatment methods as the environment created is sterile and safe. Generative models would prove to be effective for the modeling and simulation of diseases in neuroscience; the realistic simulation of disorders, tailoring treatment approaches for specific patients, and coming up with new therapeutic ways can be achieved. The breakthrough is promising to have the neurological conditions better understood and consequently improve patient care.

3.4. Brain atlas creation and individualization

Generative models are vital in the generation and customization of brain atlases, necessary for mapping the intricate complexity of the human brain in terms of its structure and function. Generative models such as VAEs, GANs, and flow-based models can be trained on large datasets of neuroimaging scans. Once the model is learned, it can generate high-resolution, subject-specific maps of the brain by processing an individual's neuroimaging data. Such subject-specific mappings include unique patterns of connectivity, cortical thickness, and subcortical structures and are very highly accurate representations of an individual's brain. These models improve the resolution and accuracy of brain atlases by integrating high-dimensional data. This allows for the development of comprehensive atlases capturing both structural and functional features of the brain with high precision. Therefore, personalized atlases of the brain are likely to provide more detailed and nuanced insights into the organization of the brain, an important aspect of research and clinical applications. Personalized atlases of the brain are of great clinical utility. These models can improve the diagnosis and treatment of neurological disorders because they provide a detailed mapping of an individual's brain, which would help one identify abnormal regions and then guide interventions, such as surgical planning or targeted therapies. Before deciding where on the body to apply force, neurosurgical robots can construct a detailed and personal topographic map of the brain, so that dangerous zones can be avoided during surgery. In research, localized brain atlases are more truthful in measuring the brain's behavior and the impact of distinct neurological disorders. Researchers can then use these atlases in search of how the given difference in connectivity and brain anatomy between one individual and the other is linked to a specific cognitive function or behavioral outcome in the working of the brain. The new generative models will be changing the generation and personalization of brain atlases by increasing their accuracy, resolution, and personalization. Such designs usher in the progressive potential for enhancing the capacity to advance neuroscience as a tool for enhancing the course of therapy and research studies.

3.5. Anomaly detection

Neuroimaging is a major area that has benefited from generative models in the case of anomaly detection. These are capable of identifying aberrations that may represent disease or injury and subsequently enhance the results of diagnosis by offering prospects for timely treatment. Some of the advantages of generative models, including VAEs and

GANs, are: the ability to capture the distribution of neuroimaging data. If these models had been trained on large datasets of healthy brains, they would have encoded the normal variability in the brain and these structures. After it has been trained, the models can then be used to compare new neuroimaging data to this learned distribution for abnormalities or shifts. For instance, if you carry out a generative model with MRI scans, it is possible to point out certain regions as being anomalous. Discrepancies could be seen in terms of structural similarity, different patterns of growth, or an odd rhythm that could all be signs of illnesses, for example, tumors, Alzheimer's disease, or multiple sclerosis. Such anomalies are reported to radiologists and neurologists so that attention will be shifted to areas that are considered to be problematic and require more studies. Generative models can also identify movements that the naked eye is unable to spot, or at the very least, does not notice. Such models can be used in the analysis of functional MRI data for the detection of abnormally structured signals, which can be related primarily to some diseases, for example, epilepsy or traumatic brain injury. Generative models add a depth of diagnostic support for how much a certain scan abstracts from the norm. The specific capacity of flow-based models for computing the exact likelihood is especially valuable for the detection of anomalies. These models can make a probability assignment of each data point, and the exact localization of abnormalities can be defined. This capability is extremely useful for monitoring disease advancement and for the identification of diseases in which early intervention can result in a favorable patient prognosis. Generative models are useful in the augmentation of neuroimaging with the learning of normal brains and identifying alterations that present as disease or injury. Their potential to provide detailed and measurable indications of any abnormalities makes them indispensable when it comes to the correction of diagnostic blunders and early identification of neurological disorders.

3.6. Personalized medicine

Generative models have had a revolutionary intervention on personal medicine, most especially in neuroimaging and neurological disorders. Personalized Diagnosis: Extended generative models can take a patient's neuroimaging information to determine specific patterns and anomalies that may not be apparent by other means of analysis. These models produce pseudo-images to represent the patients' brain performance, thus aiding in the development of a fairly accurate and unique diagnostic outlook. Such an approach in diagnosing helps identify neurological disorders, for instance, Alzheimer's disease, multiple sclerosis, or brain tumors in their early stages. This may also significantly increase the probability of a correct diagnosis. As soon as one is diagnosed with a condition, generative models can be used to come up with treatment plans. By modeling the influence of the treatments on an individual's brain, the doctors can estimate how a certain patient would react in the presence of alternative fine treatment methods. It aids in treatment plan choice, side effect reduction, and enhancing patients' status. For instance, with generative models, one can think of an application in the ability to handle epilepsy or the prognosis of different possibilities of surgical operations that should spare critical areas in the patient's brain for surgeons. Generative models also help in personalized prognosis and predict the progression of neurological diseases by analyzing longitudinal neuroimaging data. These models can generate future scenarios for disease progression. This predictive capability is critical when managing chronic neurological conditions, especially in continuous monitoring and timely intervention-where quality of life can substantially be improved. Generative models have the capability to develop virtual representations of patients' brains, replicating their distinct attributes and functionality. These digital replicas provide a safe environment for analyzing various medical conditions, predicting disease progression, and evaluating potential treatments without any risk to real patients. This innovative approach enhances our comprehension of individual differences in neurological disorders and accelerates the creation of targeted therapies tailored to each patient's specific needs. By leveraging generative models, researchers can significantly improve neuroimaging quality, personalize medical care, and refine diagnostic and treatment methodologies. The practical applications of generative models, as outlined in Table 6, showcase their extensive impact on neuroimaging. These models are not only valuable in refining image quality but also play a crucial role in improving diagnostic precision, predicting disease outcomes, and customizing treatment strategies. Various approaches, including GANs, VAEs, and digital twin technology, are helping to develop more reliable and adaptable AI systems. The integration of these methods with traditional machine learning techniques presents new possibilities for enhancing clinical diagnostics and medical decision-making.

Advanced techniques such as GANs and VAEs contribute to expanding neuroimaging datasets by generating synthetic images, thereby improving model accuracy and robustness. Convolutional autoencoders (CAEs) and GANs help refine image quality by removing noise, enhancing resolution, and eliminating distortions. Super-Resolution GANs (SRGANs) improve image clarity for better anatomical studies, while conditional GANs facilitate cross-modal imaging, generating missing scan types as needed. Autoencoder-based models are instrumental in identifying structural abnormalities in the brain, thereby aiding in the detection of neurological diseases. Diffusion models further enhance neuroimaging by filling in missing data, leading to more comprehensive imaging results. The combination of digital twin technology, VAEs, and deep learning supports the simulation of disease progression and treatment responses, fostering the development of personalized medical interventions. Additionally, generative models enable the creation of detailed, patient-specific brain atlases that reflect unique anatomical and functional characteristics, thereby improving accuracy in both research and clinical applications. These AI-driven approaches have proven effective in detecting neurological disorders such as Alzheimer's disease and brain tumors with greater precision. Through personalized simulations and digital twin models, generative AI aids in formulating customized treatment plans, ultimately improving patient outcomes.

4. Challenges and future directions

Despite their transformative potential, generative AI applications in neuroimaging encounter multiple challenges, including data limitations, interpretability issues, computational resource constraints, generalization difficulties, and ethical concerns. One of the primary obstacles is the scarcity of high-quality neuroimaging datasets, as acquiring such data is expensive and subject to strict ethical and privacy regulations. The inherent variability in imaging modalities, such as MRI, fMRI, and PET scans, further complicates model training, making it difficult to develop AI systems that perform consistently across different datasets. Demographic imbalances in training data can introduce biases, leading to reduced model performance for underrepresented groups. To address these issues, researchers must focus on improving data augmentation techniques, standardizing imaging protocols, and developing bias mitigation strategies to enhance fairness and reliability. Other challenges include instability in GAN training, high computational costs, and difficulties in interpreting model decisions, which hinder their adoption in clinical settings. Overcoming these limitations requires the development of stable model architectures, efficient training methodologies, and domain-specific optimizations. Generative models hold immense potential to advance neuroimaging by refining image quality, improving diagnostics, and enabling precision treatment. Their applications extend beyond imaging enhancements to include brain mapping, disease modeling, and predictive analytics, thereby shaping the future of neuroscience research and clinical care. Continuous innovation in this field will be key to unlocking new insights into complex neurological conditions and improving patient care through AI-driven precision medicine.

Model interpretability remains a significant challenge, as the complexity of generative AI often obscures its decision-making processes. For successful integration into clinical practice, researchers and healthcare professionals must be able to understand and validate AI-driven conclusions. The development of transparent AI architectures and user-friendly interfaces that visualize decision-making processes will be essential for building trust and facilitating clinical adoption. Enhancing interpretability will also aid in identifying potential errors and validating model outputs in real-world medical applications. The computational demands of generative models pose another major challenge, especially for institutions with limited access to high-performance computing resources. Training deep generative AI systems requires significant computational power, leading to high costs, extended processing times, and considerable energy consumption. The environmental impact of energy-intensive computations further raises sustainability concerns. To address these challenges, researchers should explore efficient model architectures, optimize hardware usage, and implement alternative computing solutions such as distributed and cloud-based processing. Reducing computational costs will be critical in making generative AI more accessible for widespread use in neuroimaging and medical applications.

Ensuring the robustness and adaptability of generative models is another critical concern, as models trained on specific datasets may not generalize well to unseen or diverse real-world data. Improving the generalization capabilities of these AI systems will enhance their clinical reliability and effectiveness in medical decision-making. Researchers must focus on advancing training methodologies and refining validation techniques to ensure generative models can adapt to diverse and noisy clinical datasets.

Ethical considerations play a crucial role in the responsible development and deployment of generative AI in medical imaging. As synthetic data becomes more prevalent in research and clinical settings, it is imperative to address privacy risks, bias mitigation, and ethical use. A major concern is data privacy, as AI models often rely on sensitive medical datasets for training. This raises the risk of data breaches, patient re-identification, and confidentiality violations. To safeguard patient privacy, robust governance frameworks must be established, incorporating measures such as differential privacy, data anonymization, and secure training protocols. However, achieving complete anonymization remains a challenge, as sophisticated re-identification techniques can sometimes extract personal data from seemingly anonymous datasets. Balancing privacy protection with AI model effectiveness is a key consideration for ensuring secure and ethical AI applications. Bias in AI-generated neuroimaging data is another pressing issue, as models can inherit and amplify biases present in training datasets. These biases can lead to disparities in diagnostic accuracy and treatment recommendations, particularly for underrepresented populations. To address this, researchers must implement fairness-aware algorithms, adversarial debiasing techniques, and diverse training datasets to improve model equity and reliability. Ensuring demographic inclusivity in training data is essential to preventing biased AI outcomes and promoting fairer healthcare applications.

Another ethical risk involves the potential misuse of synthetic data in medical applications. While generative AI offers numerous benefits, it also presents the possibility of fraudulent medical records, manipulated diagnostic data, and the creation of deepfake medical images. Such unethical uses could erode public trust in AI-driven healthcare and cause significant harm. In research settings, falsified synthetic data could distort scientific findings, leading to flawed medical conclusions. Similarly, in commercial applications, unethical use of AI-generated medical data could lead to misleading marketing or exploitative practices. To prevent such risks, regulatory frameworks and ethical guidelines must be established to ensure responsible AI deployment in neuroimaging. Transparency in model development and decision-making will be crucial in maintaining accountability and trust in AI applications.

Addressing the computational and ethical challenges associated with generative AI requires continuous advancements in privacy-preserving techniques, fairness optimization, and sustainable computing solutions. Efficient model architectures, transfer learning, and energy-efficient hardware can help reduce computational costs while maintaining high performance. Cloud computing and distributed processing offer scalable solutions for handling large neuroimaging datasets while minimizing environmental impact. Ensuring equitable access to AI resources is also critical in closing the gap between well-funded research institutions and under-resourced healthcare providers. A concerted effort from researchers, policymakers, and industry stakeholders will be necessary to ensure the ethical, fair, and sustainable development of generative AI in neuroimaging.

Table 7 summarizes the key challenges and future directions in applying generative models to neuroimaging, highlighting the ongoing efforts required to optimize these technologies for widespread clinical use.

Table 7: Challenges and future directions in generative models for neuroimaging

Challenges	Description	Future directions
Data limitations	Data scarcity, heterogeneity across modalities (e.g., MRI, fMRI, PET), and biases in datasets hinder the development of robust models. Limited dataset size and ethical concerns further complicate model training	Develop techniques for augmenting and harmonizing datasets. Focus on overcoming biases and increasing data diversity for better model generalization
Interpretability and explainability	Complex models are difficult to interpret, limiting their clinical adoption. Clinicians require transparency to trust and utilize these models effectively	Focus on developing interpretable and transparent models, with tools that allow clinicians to understand the decision-making process and model outputs
Computational cost	Training generative models requires high computational resources, limiting accessibility for many research institutions and healthcare providers	Design more efficient algorithms and leverage hardware optimization to reduce the computational burden, making models more accessible for practical use
Generalizability and robustness	Models trained on specific datasets may not perform well on diverse or noisy data, reducing their reliability in real-world applications	Improve model generalization and robustness, ensuring reliable performance on unseen data and in varying conditions, particularly in clinical settings
Ethical considerations	Data privacy and misuse of generated images (e.g., fake medical records) pose significant risks. Secure storage, processing, and ethical standards for model deployment are essential	Establish robust data governance frameworks and ethical guidelines to ensure patient confidentiality, transparency in model development, and ethical use of generated images
Exploring novel applications	Generative models can extend beyond traditional tasks to simulate disease progression, predict treatment outcomes, and personalize medicine	Explore new applications such as simulating disease progression, creating personalized brain models, and predicting treatment outcomes for improved healthcare delivery
Integrating multimodal data	Combining different neuroimaging modalities (e.g., MRI, fMRI, PET) can provide a comprehensive view of brain structure and function, enhancing the capabilities of generative models	Focus on integrating multimodal data to provide a more holistic understanding of brain health, improving the overall effectiveness and precision of generative models

There are several key research avenues that can enhance the effectiveness and usability of generative AI in neuroscience imaging. Enhancing computational efficiency and improving model transparency will aid in making better clinical decisions. Expanding the role of generative AI beyond standard image processing—such as modeling disease progression, forecasting treatment responses, and creating individualized brain models—can drive advancements in precision medicine. Strengthening data augmentation strategies will help mitigate data shortages while promoting more diverse and representative training datasets. Moreover, incorporating multimodal neuroimaging data by integrating MRI, fMRI, and PET scan insights can offer a more holistic view of brain structure and function, leading to more accurate and reliable AI models. Although generative AI has the potential to revolutionize neuroscience imaging, it is essential to tackle challenges such as data scarcity, computational demands, model interpretability, generalization, and ethical considerations. Addressing these hurdles will facilitate the development of more reliable, transparent, and clinically relevant AI-powered neuroimaging solutions.

5. Conclusion

This review exemplifies the transformative potential of generative AI for neuroscience imaging, from GANs, VAEs, diffusion, and flow-based models. It improves the image quality and can also augment synthetic data to overcome scarcity issues, and it also improves diagnostic accuracy by removing noise and artifacts. More importantly, generative models are essential for doing robust data augmentation, which would improve training as well as the generalization ability of the machine learning algorithm for better diagnostics. These models further extend image enhancement capabilities, such as denoising and super-resolution, to support more accurate, clearer neuroimaging critical for clinical care. Their utility in modeling and simulating disease helps reveal mechanisms, thereby accelerating personalized treatment development. Limitations in data (lucidity, heterogeneity, bias), model interpretability, computational costs, and generalization to diverse data remain open areas of future work. The development of more efficient and interpretable generative models is needed to achieve higher clinical impact. Exploring novel applications—simulating disease progression, predicting treatment outcomes, creating personalized brain atlases—holds immense potential for revolutionizing personalized medicine. Addressing ethical considerations, including data privacy and responsible use, is paramount. Generative AI promises to revolutionize neuroscience imaging, enhancing data quality, improving diagnostics, and enabling personalized medicine. While challenges remain, ongoing research will overcome these obstacles, establishing generative AI as a central tool for advancing neuroscience and improving patient care. Future research efforts should focus on creating more efficient, interpretable, and robust models, exploring novel applications, and ensuring ethical and fair use of these powerful tools. The future implications for innovation and impact are vast.

List of abbreviations

AD	Alzheimer's disease
AI	Artificial intelligence
CAEs	Convolutional autoencoders
CNN	Convolutional neural network
CT	Computed tomography
DCGAN	Deep convolutional generative adversarial network
DMCA-GAN	Dual multilevel constrained attention GAN
DTI	Diffusion tensor imaging
DVAE	Deep variational autoencoder
EEG	Electroencephalography
FC	Functional connectivity
FBN	Functional brain network
FID	Fréchet inception distance
fMRI	Functional magnetic resonance imaging
GANs	Generative adversarial networks
KL	Kullback–Leibler (divergence)
MDD	Major depressive disorder
MCI	Mild cognitive impairment
ML	Machine learning
MRI	Magnetic resonance imaging
MSE	Mean squared error
MSD	Medical Segmentation Decathlon
NMSE	Normalized mean squared error
ODF	Orientation distribution functions
PET	Positron emission tomography
PSNR	Peak signal-to-noise ratio
SFL	Split federated learning
SR	Super-resolution
SSIM	Structural similarity index
SSA-CycleGAN	Semi-supervised attention-guided CycleGAN
TES	Transcranial electrical stimulation
VAEs	Variational autoencoders
WGANs	Wasserstein GANs

Compliance with ethical standards

Conflict of interest

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

References

Abbasi S, Lan H, Choupan J, Sheikh-Bahaei N, Pandey G, and Varghese B (2024). Deep learning for the harmonization of structural MRI scans: A survey. BioMedical Engineering OnLine, 23: 90. https://doi.org/10.1186/s12938-024-01280-6 **PMid:39217355 PMCid:PMC11365220**

Ahmadi F, Shiri ME, Bidabad B, Sedaghat M, and Memari P (2024). Ricci flow-based brain surface covariance descriptors for diagnosing Alzheimer's disease. Biomedical Signal Processing and Control, 93: 106212. https://doi.org/10.1016/j.bspc.2024.106212

Bacon EJ, He D, Achi NBADA, Wang L, Li H, Yao-Digba PDZ, Monkam P, and Qi S (2024). Neuroimage analysis using artificial intelligence approaches: A systematic review. Medical and Biological Engineering and Computing, 62(9): 2599-2627. https://doi.org/10.1007/s11517-024-03097-w **PMid:38664348**

Bit S, Dey P, Maji A et al. (2024). MRI-based mild cognitive impairment and Alzheimer's disease classification using an algorithm of combination of variational autoencoder and other machine learning classifiers. Journal of Alzheimer's Disease Reports, 8(1): 1434-1452. https://doi.org/10.1177/25424823241290694 **PMid:40034356 PMCid:PMC11863754**

Bowles C, Chen L, Guerrero R et al. (2018). GAN augmentation: Augmenting training data using generative adversarial networks. Arxiv Preprint Arxiv:1810.10863. https://doi.org/10.48550/arXiv.1810.10863

Bui TD, Nguyen M, Le N, and Luu K (2020). Flow-based deformation guidance for unpaired multi-contrast MRI image-to-image translation. In the 23rd International Conference on Medical Image Computing and Computer Assisted Intervention (Part II), Springer International Publishing, Lima, Peru: 728-737. https://doi.org/10.1007/978-3-030-59713-9_70

Cao H, Tan C, Gao Z, Xu Y, Chen G, Heng PA, and Li SZ (2024). A survey on generative diffusion models. IEEE Transactions on Knowledge and Data Engineering, 36(7): 2814-2830. https://doi.org/10.1109/TKDE.2024.3361474

Chadebec C, Thibeau-Sutre E, Burgos N, and Allassonnière S (2022). Data augmentation in high dimensional low sample size setting using a geometry-based variational autoencoder. IEEE Transactions on Pattern Analysis and Machine Intelligence, 45(3): 2879-2896. https://doi.org/10.1109/TPAMI.2022.3185773 **PMid:35749321**

Chatterjee S, Sciarra A, Dünnwald M et al. (2022). StRegA: Unsupervised anomaly detection in brain MRIs using a compact context-encoding variational autoencoder. Computers in Biology and Medicine, 149: 106093. https://doi.org/10.1016/j.compbiomed.2022.106093 **PMid:36116318**

Chen X, Peng Y, Li D, and Sun J (2023). DMCA-GAN: Dual multilevel constrained attention GAN for MRI-based hippocampus segmentation. Journal of Digital Imaging, 36(6): 2532-2553. https://doi.org/10.1007/s10278-023-00854-5 **PMid:37735310 PMCid:PMC10584805**

Choi KS and Sunwoo L (2022). Artificial intelligence in neuroimaging: Clinical applications. Investigative Magnetic Resonance Imaging, 26(1): 1-9. https://doi.org/10.13104/imri.2022.26.1.1

Fan Y, Liao H, Huang S, Luo Y, Fu H, and Qi H (2024). A survey of emerging applications of diffusion probabilistic models in MRI. Meta-Radiology, 2(2): 100082. https://doi.org/10.1016/j.metrad.2024.100082

Gajjar P, Garg M, Desai S, Chhinkaniwala H, Sanghvi HA, Patel RH, Gupta S, and Pandya AS (2024). An empirical analysis of diffusion, autoencoders, and adversarial deep learning models for predicting dementia using high-fidelity MRI. IEEE Access, 12: 131231-131243. https://doi.org/10.1109/ACCESS.2024.3354724

Gong C, Jing C, Chen X et al. (2023). Generative AI for brain image computing and brain network computing: A review. Frontiers in Neuroscience, 17: 1203104. https://doi.org/10.3389/fnins.2023.1203104 **PMid:37383107 PMCid:PMC10293625**

Goodfellow IJ, Pouget-Abadie J, Mirza M et al. (2014). Generative adversarial nets. Arxiv Preprint Arxiv:1406.2661v1. https://doi.org/10.48550/arXiv.1406.2661

Gupta P, Ding B, Guan C, and Ding D (2024). Generative AI: A systematic review using topic modelling techniques. Data and Information Management, 8(2): 100066. https://doi.org/10.1016/j.dim.2024.100066

Gupta R, Tiwari S, and Chaudhary P (2025). Large generative models for different data types. In: Gupta R, Tiwari S, and Chaudhary P (Eds.), Generative AI: Techniques, models and applications: 103-162. Springer, Cham, Switzerland. https://doi.org/10.1007/978-3-031-82062-5_6

Hagos DH, Battle R, and Rawat DB (2024). Recent advances in generative AI and large language models: Current status, challenges, and perspectives. IEEE Transactions on Artificial Intelligence, 5(12): 5873-5893. https://doi.org/10.1109/TAI.2024.3444742

Han C, Rundo L, Araki R, Nagano Y, Furukawa Y, Mauri G, Nakayama H, and Hayashi H (2019). Combining noise-to-image and image-to-image GANs: Brain MR image augmentation for tumor detection. IEEE Access, 7: 156966-156977. https://doi.org/10.1109/ACCESS.2019.2947606

Hwang S and Shin J (2024). Prognosis prediction of Alzheimer's disease based on multi-modal diffusion model. In the 18th International Conference on Ubiquitous Information Management and Communication, IEEE, Kuala Lumpur, Malaysia: 1-5. https://doi.org/10.1109/IMCOM60618.2024.10418423

Jeevan P, Nixon N, and Sethi A (2024). Normalizing flow-based metric for image generation. Arxiv Preprint Arxiv:2410.02004. https://doi.org/10.48550/arXiv.2410.02004

Ledig C, Theis L, Huszár F, Caballero J, Cunningham A, Acosta A, Aitken A, Tejani A, Totz J, Wang Z, and Shi W (2017). Photo-realistic single image super-resolution using a generative adversarial network. In the IEEE Conference on Computer Vision and Pattern Recognition, IEEE, Honolulu, USA: 4681-4690. https://doi.org/10.1109/CVPR.2017.19

Molnár S and Tamás L (2024). Variational autoencoders for 3D data processing. Artificial Intelligence Review, 57: 42. https://doi.org/10.1007/s10462-023-10687-x

Nalepa J, Marcinkiewicz M, and Kawulok M (2019). Data augmentation for brain-tumor segmentation: A review. Frontiers in Computational Neuroscience, 13: 83. https://doi.org/10.3389/fncom.2019.00083 **PMid:31920608 PMCid:PMC6917660**

Nie D, Trullo R, Lian J, Petitjean C, Ruan S, Wang Q, and Shen D (2017). Medical image synthesis with context-aware generative adversarial networks. In the 20th International Conference on Medical Image Computing and Computer Assisted Intervention (Part III), Springer International Publishing, Quebec, Canada: 417-425. https://doi.org/10.1007/978-3-319-66179-7_48 **PMid:30009283 PMCid:PMC6044459**

Nimeshika GN and Subitha D (2024). Enhancing Alzheimer's disease classification through split federated learning and GANs for imbalanced datasets. PeerJ Computer Science, 10: e2459. https://doi.org/10.7717/peerj-cs.2459 **PMid:39650412 PMCid:PMC11623002**

O'Donnell LJ and Schultz T (2015). Statistical and machine learning methods for neuroimaging: Examples, challenges, and extensions to diffusion imaging data. In: Hotz I and Schultz T (Eds.), Visualization and processing of higher order descriptors for multi-valued data: Mathematics and visualization: 299-319. Springer, Cham, Switzerland. https://doi.org/10.1007/978-3-319-15090-1_15

Qiang N, Dong Q, Ge F, Liang H, Ge B, Zhang S, Sun Y, Gao J, and Liu T (2020). Deep variational autoencoder for mapping functional brain networks. IEEE Transactions on Cognitive and Developmental Systems, 13(4): 841-852. https://doi.org/10.1109/TCDS.2020.3025137

Rais K, Amroune M, Benmachiche A, and Haouam MY (2024). Exploring variational autoencoders for medical image generation: A comprehensive study. Arxiv Preprint Arxiv:2411.07348. https://doi.org/10.48550/arXiv.2411.07348

Reaungamornrat S, Sari H, Catana C, and Kamen A (2022). Multimodal image synthesis based on disentanglement representations of anatomical and modality specific features, learned using uncooperative relativistic GAN. Medical Image Analysis, 80: 102514. https://doi.org/10.1016/j.media.2022.102514 **PMid:35717874 PMCid:PMC9810205**

Sabuhi M, Zhou M, Bezemer CP, and Musilek P (2021). Applications of generative adversarial networks in anomaly detection: A systematic literature review. IEEE Access, 9: 161003-161029. https://doi.org/10.1109/ACCESS.2021.3131949

Sajjad M, Ramzan F, Khan MUG, Rehman A, Kolivand M, Fati SM, and Bahaj SA (2021). Deep convolutional generative adversarial network for Alzheimer's disease classification using positron emission tomography (PET) and synthetic data augmentation. Microscopy Research and Technique, 84(12): 3023-3034. https://doi.org/10.1002/jemt.23861 **PMid:34245203**

Seiler M and Ritter K (2025). Pioneering new paths: The role of generative modelling in neurological disease research. Pflügers Archiv-European Journal of Physiology, 477: 571-589. https://doi.org/10.1007/s00424-024-03016-w **PMid:39377960 PMCid:PMC11958445**

Shin HC, Tenenholtz NA, Rogers JK, Schwarz CG, Senjem ML, Gunter JL, Andriole KP, and Michalski M (2018). Medical image synthesis for data augmentation and anonymization using generative adversarial networks. In the 3rd International Workshop on Simulation and Synthesis in Medical Imaging, Springer International Publishing, Granada, Spain: 1-11. https://doi.org/10.1007/978-3-030-00536-8_1 **PMid:37312871**

Sidulova M and Park CH (2023). Conditional variational autoencoder for functional connectivity analysis of autism spectrum disorder functional magnetic resonance imaging data: A comparative study. Bioengineering, 10(10): 1209.

https://doi.org/10.3390/bioengineering10101209
PMid:37892939 PMCid:PMC10604768

Singh A and Ogunfunmi T (2021). An overview of variational autoencoders for source separation, finance, and bio-signal applications. Entropy, 24(1): 55.
https://doi.org/10.3390/e24010055
PMid:35052081 PMCid:PMC8774760

Song W, Wang X, Jiang Y, Li S, Hao A, Hou X, and Qin H (2024). Expressive 3D facial animation generation based on local-to-global latent diffusion. IEEE Transactions on Visualization and Computer Graphics, 30(11): 7397-7407.
https://doi.org/10.1109/TVCG.2024.3456213
PMid:39255115

Sun H, Mehta R, Zhou HH, Huang Z, Johnson SC, Prabhakaran V, and Singh V (2019). Dual-glow: Conditional flow-based generative model for modality transfer. In the IEEE/CVF International Conference on Computer Vision, IEEE, Seoul, Korea: 10611-10620.
https://doi.org/10.1109/ICCV.2019.01071
PMid:35125977 PMCid:PMC8813086

Tan YF, Ting CM, Noman F, Phan RCW, and Ombao H (2024). fMRI functional connectivity augmentation using convolutional generative adversarial networks for brain disorder classification. In the IEEE International Symposium on Biomedical Imaging, IEEE, Athens, Greece: 1-5.
https://doi.org/10.1109/ISBI56570.2024.10635290

Tomczak JM (2020). General invertible transformations for flow-based generative modeling. Arxiv Preprint Arxiv:2011.15056.
https://doi.org/10.48550/arXiv.2011.15056

Wang D, Ma C, and Sun S (2023a). Novel paintings from the latent diffusion model through transfer learning. Applied Sciences, 13(18): 10379. https://doi.org/10.3390/app131810379

Wang R, Bashyam V, Yang Z et al. (2023b). Applications of generative adversarial networks in neuroimaging and clinical neuroscience. Neuroimage, 269: 119898.
https://doi.org/10.1016/j.neuroimage.2023.119898
PMid:36702211 PMCid:PMC9992336

Wang Y, Solera-Rico A, Vila CS, and Vinuesa R (2024a). Towards optimal β-variational autoencoders combined with transformers for reduced-order modelling of turbulent flows. International Journal of Heat and Fluid Flow, 105: 109254.
https://doi.org/10.1016/j.ijheatfluidflow.2023.109254

Wang Z, Li D, Wu Y, He T, Bian J, and Jiang R (2024b). Diffusion models in 3D vision: A survey. Arxiv Preprint Arxiv:2410.04738.
https://doi.org/10.48550/arXiv.2410.04738

Wei B, Wen Y, Liu X, Qi X, and Sheng B (2023). SOFNet: Optical-flow based large-scale slice augmentation of brain MRI. Displays, 80: 102536.
https://doi.org/10.1016/j.displa.2023.102536

Wei R and Mahmood A (2020). Recent advances in variational autoencoders with representation learning for biomedical informatics: A survey. IEEE Access, 9: 4939-4956.
https://doi.org/10.1109/ACCESS.2020.3048309

Xu Z, Qi C, and Xu G (2019). Semi-supervised attention-guided cyclegan for data augmentation on medical images. In the IEEE International Conference on Bioinformatics and Biomedicine, IEEE, San Diego, USA: 563-568.
https://doi.org/10.1109/BIBM47256.2019.8982932

Yen C, Lin CL, and Chiang MC (2023). Exploring the frontiers of neuroimaging: A review of recent advances in understanding brain functioning and disorders. Life, 13(7): 1472.
https://doi.org/10.3390/life13071472
PMid:37511847 PMCid:PMC10381462

Yilmaz B and Korn R (2024). A comprehensive guide to generative adversarial networks (GANs) and application to individual electricity demand. Expert Systems with Applications, 250: 123851. https://doi.org/10.1016/j.eswa.2024.123851

Yin J, Qiao Z, Han L, and Zhang X (2025). EEG-based emotion recognition with autoencoder feature fusion and MSC-TimesNet model. Computer Methods in Biomechanics and Biomedical Engineering.
https://doi.org/10.1080/10255842.2025.2477801

Yoon D, Myong Y, Kim YG, Sim Y, Cho M, Oh BM, and Kim S (2024). Latent diffusion model-based MRI superresolution enhances mild cognitive impairment prognostication and Alzheimer's disease classification. NeuroImage, 296: 120663.
https://doi.org/10.1016/j.neuroimage.2024.120663
PMid:38843963

Yuan C, Duan J, Xu K, Tustison NJ, Hubbard RA, and Linn KA (2024). ReMiND: Recovery of missing neuroimaging using diffusion models with application to Alzheimer's disease. Imaging Neuroscience, 2: 1-14.
https://doi.org/10.1162/imag_a_00323
PMCid:PMC12290738

Zhen X, Chakraborty R, Yang L, and Singh V (2021). Flow-based generative models for learning manifold to manifold mappings. In the Proceedings of the AAAI Conference on Artificial Intelligence, 35(12): 11042-11052.
https://doi.org/10.1609/aaai.v35i12.17318 **PMid:34457995**

Zhou T, Liu M, Thung KH, and Shen D (2019). Latent representation learning for Alzheimer's disease diagnosis with incomplete multi-modality neuroimaging and genetic data. IEEE Transactions on Medical Imaging, 38(10): 2411-2422.
https://doi.org/10.1109/TMI.2019.2913158
PMid:31021792 PMCid:PMC8034601