

PAPER • OPEN ACCESS

Beware of diffusion models for synthesizing medical images—a comparison with GANs in terms of memorizing brain MRI and chest x-ray images

To cite this article: Muhammad Usman Akbar *et al* 2025 *Mach. Learn.: Sci. Technol.* **6** 015022

View the [article online](#) for updates and enhancements.

You may also like

- [An efficient Wasserstein-distance approach for reconstructing jump-diffusion processes using parameterized neural networks](#)
Mingtao Xia, Xiangting Li, Qijing Shen *et al.*
- [Refinable modeling for unbinned SMEFT analyses](#)
Robert Schöfbeck
- [Paired autoencoders for likelihood-free estimation in inverse problems](#)
Matthias Chung, Emma Hart, Julianne Chung *et al.*



PAPER

OPEN ACCESS

RECEIVED
11 January 2024REVISED
25 November 2024ACCEPTED FOR PUBLICATION
3 December 2024PUBLISHED
29 January 2025

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Beware of diffusion models for synthesizing medical images—a comparison with GANs in terms of memorizing brain MRI and chest x-ray images

Muhammad Usman Akbar^{1,3,4} , Wuhao Wang² and Anders Eklund^{1,2,3,*} ¹ Department of Biomedical Engineering, Linköping University, Linköping, Sweden² Division of Statistics & Machine Learning, Department of Computer and Information Science, Linköping University, Linköping, Sweden³ Center for Medical Image Science and Visualization (CMIV), Linköping University, Linköping, Sweden⁴ Clinical Memory Research Unit, Department of Clinical Sciences, Lund University, Malmö, Sweden

* Author to whom any correspondence should be addressed.

E-mail: anders.eklund@liu.se**Keywords:** brain MRI, CXR, generative AI, GANs, diffusion model, synthetic data, memorization

Abstract

Diffusion models were initially developed for text-to-image generation and are now being utilized to generate high quality synthetic images. Preceded by generative adversarial networks (GANs), diffusion models have shown impressive results using various evaluation metrics. However, commonly used metrics such as Frechet inception distance and inception score are not suitable for determining whether diffusion models are simply reproducing the training images. Here we train StyleGAN and a diffusion model, using BRATS20, BRATS21 and a chest x-ray (CXR) pneumonia dataset, to synthesize brain MRI and CXR images, and measure the correlation between the synthetic images and all training images. Our results show that diffusion models are more likely to memorize the training images, compared to StyleGAN, especially for small datasets and when using 2D slices from 3D volumes. Researchers should be careful when using diffusion models (and to some extent GANs) for medical imaging, if the final goal is to share the synthetic images.

1. Introduction

Synthetic medical data have become a very popular research topic recently [18, 34, 37, 44], especially in the medical imaging domain where obtaining large datasets for deep learning is difficult due to different regulations such as GDPR (and anonymizing images is more difficult compared to anonymizing tabular data). Synthetic medical images have the potential to facilitate data sharing [2, 8], but discussions with legal experts require quantification of memorization and data leakage. Here we therefore investigate the degree of memorization for different generative models.

Generative adversarial networks (GANs) [19] have for almost a decade been used to generate images with higher and higher resolution. The machine learning community has recently paid a lot of attention to the powerful and adaptable framework known as diffusion models [21], which is another type of generative model inspired by anisotropic diffusion for image denoising. In the context of generative models, diffusion refers to the gradual transformation of random noise into meaningful data, such as images or text, through a series of carefully designed steps. The sampling of realistic samples from complex, high-dimensional distributions is the primary objective of both GANs and diffusion models, and diffusion models have been shown to outperform GANs [16]. In this work we focus on the topic of memorization, i.e. that generative models can simply generate samples that are copies of the training data. While the main focus is diffusion models, we also provide a comparison with GANs in order to better understand memorization.

1.1. GANs

The development of generative image models has a rich past. It was the emergence of GANs [19] that first made it possible to produce high-quality images on a large scale [11, 22]. GANs were groundbreaking in deep

Table 1. A comparison of previous papers which used diffusion models to synthesize medical images from noise. For papers which synthesize several types of images or volumes, we list one row per modality. There is a wide range in the number of subjects (or images) used for training the diffusion models. Based on our results, diffusion models trained with a small number of subjects or images are likely to memorize the training data.

Paper	Number of training subjects and/or images	Dimension	Modality	Mentions memorization
Ali <i>et al</i> [3]	3165 images	2D	chest x-ray	No
Moghadam <i>et al</i> [29]	344 (33 777 patches)	2D	histology	No
Pan <i>et al</i> [31]	5500 images	2D	chest x-ray	No
Pan <i>et al</i> [31]	101 (1902 images)	2D	heart MRI	No
Pan <i>et al</i> [31]	94 (4157 images)	2D	pelvic CT	No
Pan <i>et al</i> [31]	30 (2178 images)	2D	abdomen CT	No
Peng <i>et al</i> [32]	3000 images	2D	brain MRI	No
Txurio <i>et al</i> [43]	447 (57 216 images)	2D	pelvic CT	No
Packhäuser <i>et al</i> [30]	30 805 (112 120 images)	2D	chest x-ray	Yes
Pinaya <i>et al</i> [33]	31 740	3D	brain MRI	No
Khader <i>et al</i> [26]	998	3D	brain MRI	No
Khader <i>et al</i> [26]	1010	3D	thoracic CT	No
Khader <i>et al</i> [26]	1844	3D	breast MRI	No
Khader <i>et al</i> [26]	1250	3D	knee MRI	No
Dar <i>et al</i> [15]	65 (300 sub-volumes)	3D	coronary CTA	Yes
Dar <i>et al</i> [15]	904	3D	knee MRI	Yes

learning, and allows the creation of realistic and high-quality data, such as images, text, and sound. GANs have two parts: a generator and a discriminator. The generator makes synthetic data, and the discriminator determines if the data is real or synthetic. The generator and discriminator work in competition to each other to produce more and more realistic data. StyleGAN [22] is a special kind of GAN that is famous for making detailed and clear images, while giving great control over the final result.

1.2. Diffusion models

In the past couple of years, diffusion models [21] have primarily overtaken GANs, achieving top-notch performance in academic benchmarks [16] and serving as the foundation for widely-recognized image generators such as Stable Diffusion [38], DALL-E 2 [35, 36], Imagen [39] and others. One can say that diffusion probabilistic models are conceptually rather simple as they can be classified as image denoisers [21]. Despite their primary objective being to remove noise, diffusion models are capable of generating high-quality images from random noise. This denoising process occurs incrementally rather than in a single step. Diffusion models operate iteratively, progressively changing the noise to produce a refined image. This is achieved by training the model with varying degrees of noise, enabling it to adapt and produce high-quality results.

1.3. Evaluation metrics

How to evaluate the quality and diversity of synthetic images is an open research question and many different metrics have been proposed [9, 10]. Borji [9] review and critically discuss more than 24 quantitative and 5 qualitative metrics, including the most common metrics Frechet inception distance (FID) [20] and inception score (IS) [40]. Another kind of evaluation is to train a classifier or a segmentation network using the synthetic images [2, 17], and then test the networks on real images. This approach will better reflect how useful the synthetic images are, but it can be very time consuming if many generative models have been trained. Many metrics, such as FID and IS, do not consider memorization or overfitting. A possible explanation for this is that memorization has in general not been a problem for GANs.

1.4. Diffusion models in medical imaging

In medical imaging, diffusion models are still relatively rare; GANs are still the most common architecture for noise-to-image and image-to-image generation [44]. Some notable exceptions are covered here, see table 1 for an overview. Pan *et al* [31] implemented a diffusion model with transformer based denoising, instead of the common U-Net based approach, and used it to generate chest x-rays (CXRs), heart MRI, pelvic CT, and abdomen CT. For a COVID-19 classification task, the baseline network obtained an accuracy of 0.88 using a pure real dataset, 0.89 using a pure synthetic dataset, and 0.93 using a dataset mixed of real and synthetic data. Pinaya *et al* [33] implemented a conditional 3D diffusion model to synthesize brain volumes, and was able to control for example age and ventricular volume. However, the synthetic volumes were not used to train any model. Moghadam *et al* [29] used diffusion models to synthesize high quality

histopathology images of brain cancer, and demonstrated better IS and FID metrics compared to a GAN. Ali *et al* [3] used diffusion models to generate lungs x-ray and CT images, and then let radiologists evaluate the images as real or synthetic. Packhäuser *et al* [30] used diffusion models to generate chest radiographs for training thoracic abnormality classification systems, and demonstrated that diffusion models outperform GANs. Khader *et al* [26] used 3D diffusion models to synthesize brain MRI, thoracic CT, breast MRI as well as knee MRI volumes. The volumes were evaluated by radiologists, and by pre-training segmentation models. Kazerouni *et al* [24] provide a survey of diffusion models in medical imaging, but the number of papers on noise-to-image diffusion models is rather low (most are mentioned in this paper).

1.5. Memorization

Both GANs and diffusion models represent advanced generative approaches, but they differ fundamentally in their training methods and mechanisms of generating new data, which impacts how they ‘remember’ training data and why diffusion models, in particular, require careful handling. There will always be some degree of memorization, as GANs and diffusion models generate synthetic images based on the images used for training. It is important to quantify this memorization, to for example avoid generating synthetic images that are copies of the training images.

GANs are trained through a competitive process between two neural networks: a generator that creates images and a discriminator that evaluates them, through an adversarial loss function. The generator only improves its output based on feedback from the discriminator, i.e. without seeing the training images, aiming to produce images indistinguishable from real examples. The fact that the generator never sees the training images will limit memorization.

Diffusion models, on the other hand, operate through a process of adding and then iteratively removing noise from a data sample. They start by gradually adding noise to an image until the original is completely obscured, and then they learn to reverse this process to generate new images from pure noise. Because this process involves learning a very detailed noise reduction path, for example by minimizing the mean squared error between original and denoised images, it inherently incorporates detailed characteristics from the training data at each step. Thus unlike GANs, diffusion models do not only learn the general structure of the data, but can potentially learn specific instances making it more prone to learning the training data.

Surprisingly, most of the papers in table 1 do not mention memorization or overfitting. Pan *et al* [31] calculated a diversity score to obtain a measure of the diversity of the synthetic images. For each synthetic image, they calculated the structural similarity index measure (SSIM) between it and every other synthetic image to find the most similar pair. This procedure was repeated for all pairs of real images, and the distributions of nearest synthetic and real SSIM were compared using Kullback–Leibler divergence. Unfortunately the authors did not calculate SSIM between each synthetic image and all training images. To investigate memorization is even more important for larger models; their diffusion model had 1.6 billion trainable parameters (which is substantially higher compared to the GANs used for comparison). Pinaya *et al* [33] also used SSIM to investigate the diversity of the generated brain volumes, but did not calculate SSIM between real and synthetic volumes. Khader *et al* [26] also used SSIM to investigate the diversity of the synthetic volumes, but did not use SSIM to detect memorization. Packhäuser *et al* [30] guaranteed patient privacy by using additional networks to remove synthetic images that are too similar to the training images. However, it is not mentioned how many synthetic images that were seen as too similar. The authors used 112 120 chest radiographs from 30 805 patients for training, and for such a large dataset it is less likely that a diffusion model can memorize a large proportion of the images. Kazerouni *et al* [24] only mention privacy in a short paragraph in their survey. Dar *et al* [15] investigated memorization for 3D diffusion models, but did not compare with a 3D GAN and did not investigate the effect of dataset size.

1.6. Related work on memorization

Here we compare a 2D GAN and a 2D diffusion model in terms of memorization, for synthetic brain MRI and CXR images, and demonstrate that diffusion models are much more prone to memorizing the training images. This can be very problematic in medical imaging where privacy is of uttermost importance. In our previous work on training segmentation networks with synthetic images [2, 27] the obtained Dice scores were used for evaluation. When using synthetic images from diffusion models, it was observed that the Dice scores were as high as when training with real images, which made us suspicious.

This work can be seen as a continuation of several previous studies. Carlini *et al* [13] and Somepalli *et al* [41] investigated memorization in diffusion models, but their work did not focus on medical images. Pinaya *et al* [33] and Khader *et al* [26] used 3D diffusion models to synthesize brain volumes but did not address memorization. Dar *et al* [15] examined memorization in 3D diffusion models but did not compare these findings with 3D GANs. Packhäuser *et al* [30] removed overly similar synthetic images but did not report the significance of the memorization problem.

2. Methods

This section describes the two distinct datasets used (brain MRIs and CXRs), as well as the two generative models (StyleGAN and a diffusion model). Furthermore it describes different metrics and approaches for quantifying memorization. Finally, it explains how we tried different hyperparameters for the generative models, to see how they affect the degree of memorization.

2.1. Data

2.1.1. Brain tumor MRI

Deep learning can be used for many tasks in brain MRI, such as brain disease diagnosis, brain age prediction, brain anomaly detection and segmentation [14], but all tasks require large datasets for training. The MR images used for this project were downloaded from the multimodal brain tumor segmentation Challenge (BRATS) [4–7, 28]. The dataset contains MR volumes of size $240 \times 240 \times 155$, and for each patient four types of MR images are available: T1-weighted, post gadolinium contrast T1-weighted, T2-weighted, and T2 fluid attenuated inversion recovery (FLAIR). The annotations cover three parts of the brain tumor: peritumoral edema (ED), necrotic and non-enhancing tumor core (NCR/NET), and GD-enhancing tumor (ET). Two datasets were used for our experiments, BRATS20 (369 subjects) and BRATS21 (1251 subjects). The large difference in the number of subjects is excellent for testing how the memorization depends on the number of training images (which was not investigated by Carlini *et al* [13]). For each dataset, 56 subjects were reserved for testing.

All 3D volumes were split into 2D slices, as a 2D GAN and a 2D diffusion model were used (3D GANs and 3D diffusion models are not yet very common). Only slices with at least 15% pixels with an intensity of more than 50 were included in the training. This resulted in a total of 23 478 5-channel images for BRATS20, and 91 271 5-channel images for BRATS21. Each slice was zero padded from 240×240 to 256×256 pixels, as the used GAN only works for resolutions that are a power of 2, and the intensity was rescaled to 0–255. The intensities for the tumor annotations were changed from [1,2,4] to [51,102,204], such that the intensity range is more similar for the 5 channels. Figure 1 shows example images from the BRATS20 and BRATS21 datasets, displaying all four modalities and their corresponding segmentation masks.

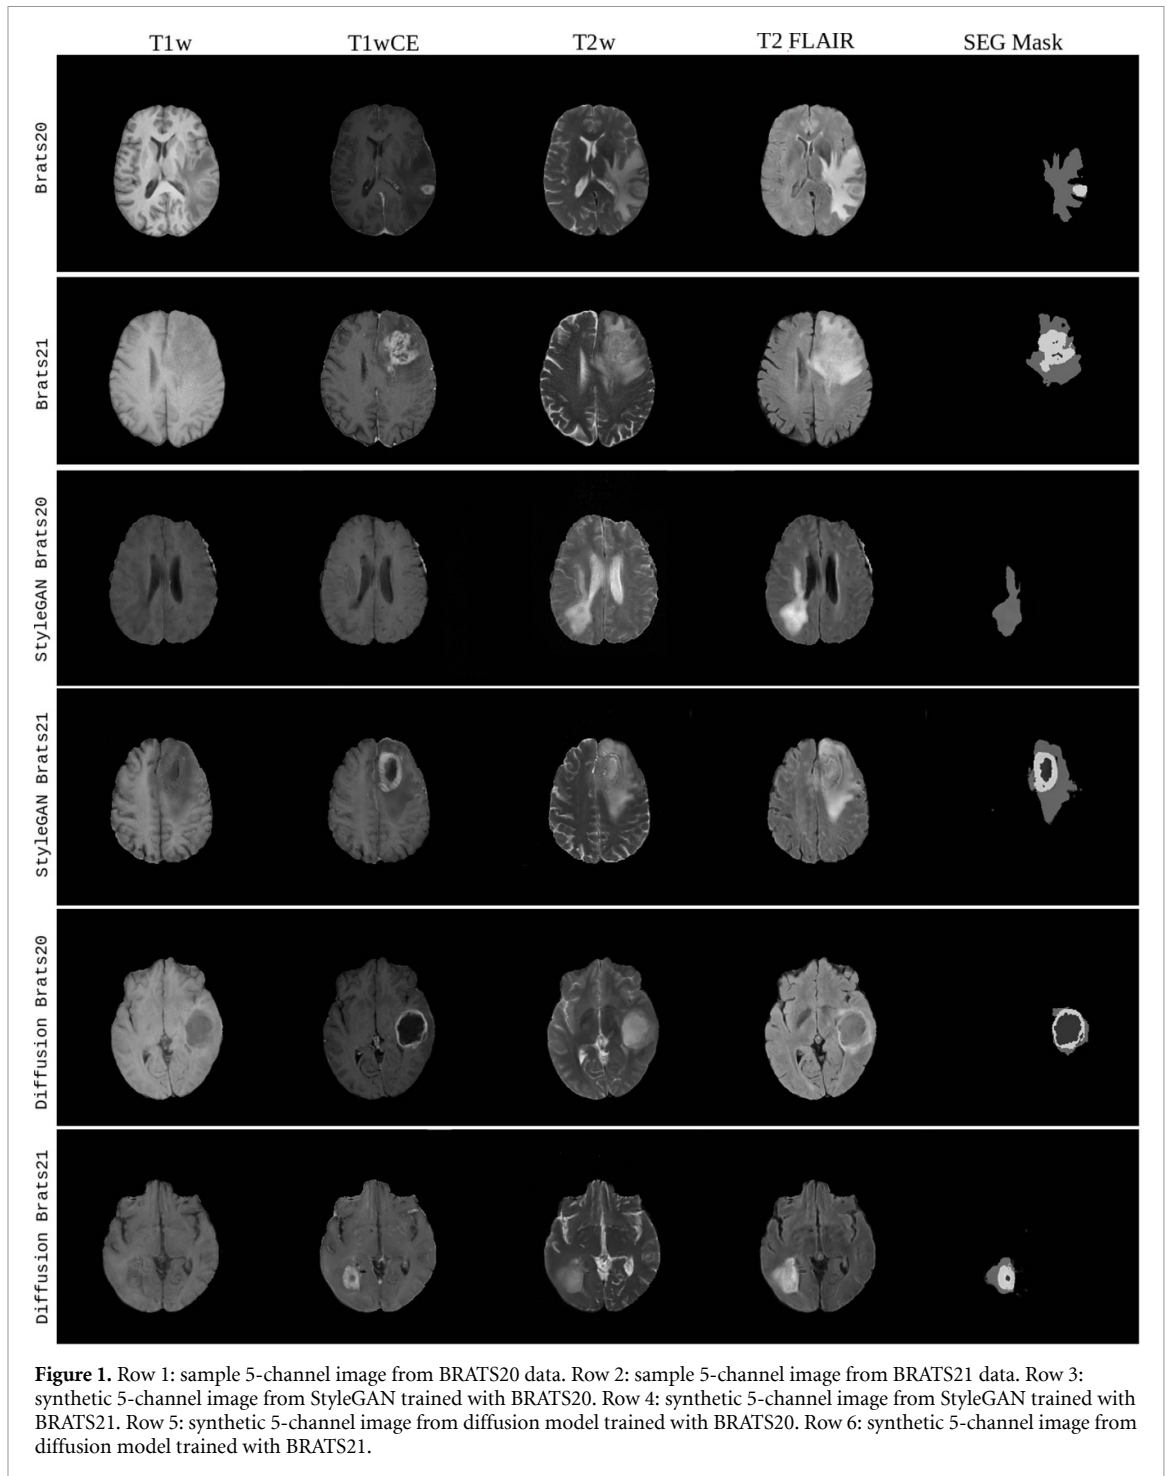
2.1.2. CXR

Deep learning can for CXRs be used for image-level prediction (classification and regression), segmentation, and localization [12], but requires large datasets for training. For comparison with 2D data, we therefore used the ‘Chest X-Ray Images (Pneumonia)’ dataset from Kaggle [25]. This dataset consists of anterior–posterior CXR images of pediatric patients aged between one to five years. For suitable alignment during GAN training, the images, which originally varied in dimensions with an average size of 1320×968 pixels, were resized to a uniform size of 256×256 pixels. The dataset comprises 5863 x-ray images in JPEG format, split into two categories: pneumonia and normal. Of these, 5216 images were allocated for training, while the remaining 624 were set aside for testing. Due to a discrepancy in the number of images between BRATS20 and BRATS21, we also employed a subset of 1300 CXR images to maintain a similar ratio. This approach facilitated a balanced comparison between CXR-5216 and CXR-1300, mirroring the comparison between BRATS20 and BRATS21.

2.2. Image generation

GANs are trained through an adversarial setup involving a generator and discriminator, where the generator aims to produce realistic images while the discriminator improves its ability to classify real versus synthetic images. This adversarial learning setup involves using an adversarial loss function. During inference, only the generator is used. On the other hand, diffusion models begin with real data samples, progressively adding noise in steps according to a linear noise schedule until the data is transformed into pure noise. The model is trained to reverse this noise process by reconstructing data from increasingly noisy states using traditional loss functions. During inference, the diffusion model starts with pure noise and iteratively applies the learned denoiser, gradually removing noise in reverse of the training noise schedule. This process results in generating an image that resembles the training data distribution.

The StyleGAN architecture [23] was used for all GAN experiments, starting with the default settings. For BRATS data, the code was modified to generate 5-channel images (T1w, T1wGD, T2w, FLAIR, segmentation annotations) instead of 3-channel color images. This enabled the GAN to learn to jointly generate 4-channel MR images and the corresponding tumor annotations simultaneously. For pneumonia data, the code was modified to generate one-channel grayscale images instead of three-channel color images. Training was



conducted with a batch size of 16, 25 000 iterations, and a gamma value of 10. The process took 3 days on 8 Nvidia A100 GPUs.

The diffusion model proposed by Ho *et al* [21] was used for all diffusion experiments, also starting with the default settings. As for StyleGAN, for BRATS data, the code was modified to generate 5-channel images instead of 3-channel color images. Training was conducted with a batch size of 12 and 7800 iterations, where each iteration corresponds to processing 1000 images. Processing was slower for diffusion models, as they perform 4000 diffusion steps for each batch. Training the diffusion model took 22 days on a single Tesla V100 GPU. For pneumonia data, the code was modified to generate one-channel grayscale images, with training taking 6 days on a single Tesla V100 GPU.

The used code for both models is publicly available on GitHub⁵

⁵ https://github.com/muhamadusman/Synthetic_Memorization

2.3. Similarity metrics for memorization

Different similarity metrics can be used to detect memorization of the training images. Image correlation is an important concept in the field of computer vision and image processing. It is a quantitative measure that gauges the similarity or relationship between two images, and can be used to study memorization. It is also an effective metric for detecting memorization in generative models because it directly compares individual pixel values between generated and training images, with high correlation indicating potential memorization. Its sensitivity to fine details allows it to capture subtle similarities, providing a clear and quantitative measure of resemblance. Additionally, pixel-wise correlation is straightforward to compute and interpret, enabling quick assessments of potential memorization without the need for complex methods.

Consider a scenario with 1000 generated images. Each generated image is compared with every original image in the training set, as it is sufficient if a single training image is memorized. For the BRATS20 dataset, which contains 23 478 training images, and the BRATS21 dataset, which contains 91 271 training images, this results in a total of $23\,478\,000$ (BRATS20) + $91\,271\,000$ (BRATS21) = $114\,749\,000$ comparisons to investigate memorization. For the CXR data, where the training set comprises 5216 and 1300 images, the total number of comparisons is $5216\,000 + 1300\,000 = 6516\,000$. This extensive comparison process ensures that even the most subtle similarities between the generated and training images are detected, providing a robust measure for identifying memorization in generative models. However, it also makes the process computationally expensive, making it challenging to add more advanced metrics due to the sheer number of comparisons required.

Other metrics, such as mutual information (MI) and structural similarity index (SSIM), can also be used to detect memorization. MI measures the amount of information shared between images, and SSIM assesses the perceived quality of images by comparing structural information. However, these metrics are more computationally demanding compared to correlation (a small comparison in Matlab revealed that SSIM is more than 10 times as demanding to calculate, and MI is 50 times as demanding as it first requires estimation of a joint histogram). Although these metrics might provide additional insights into memorization, their higher computational cost makes them less feasible for large-scale comparisons.

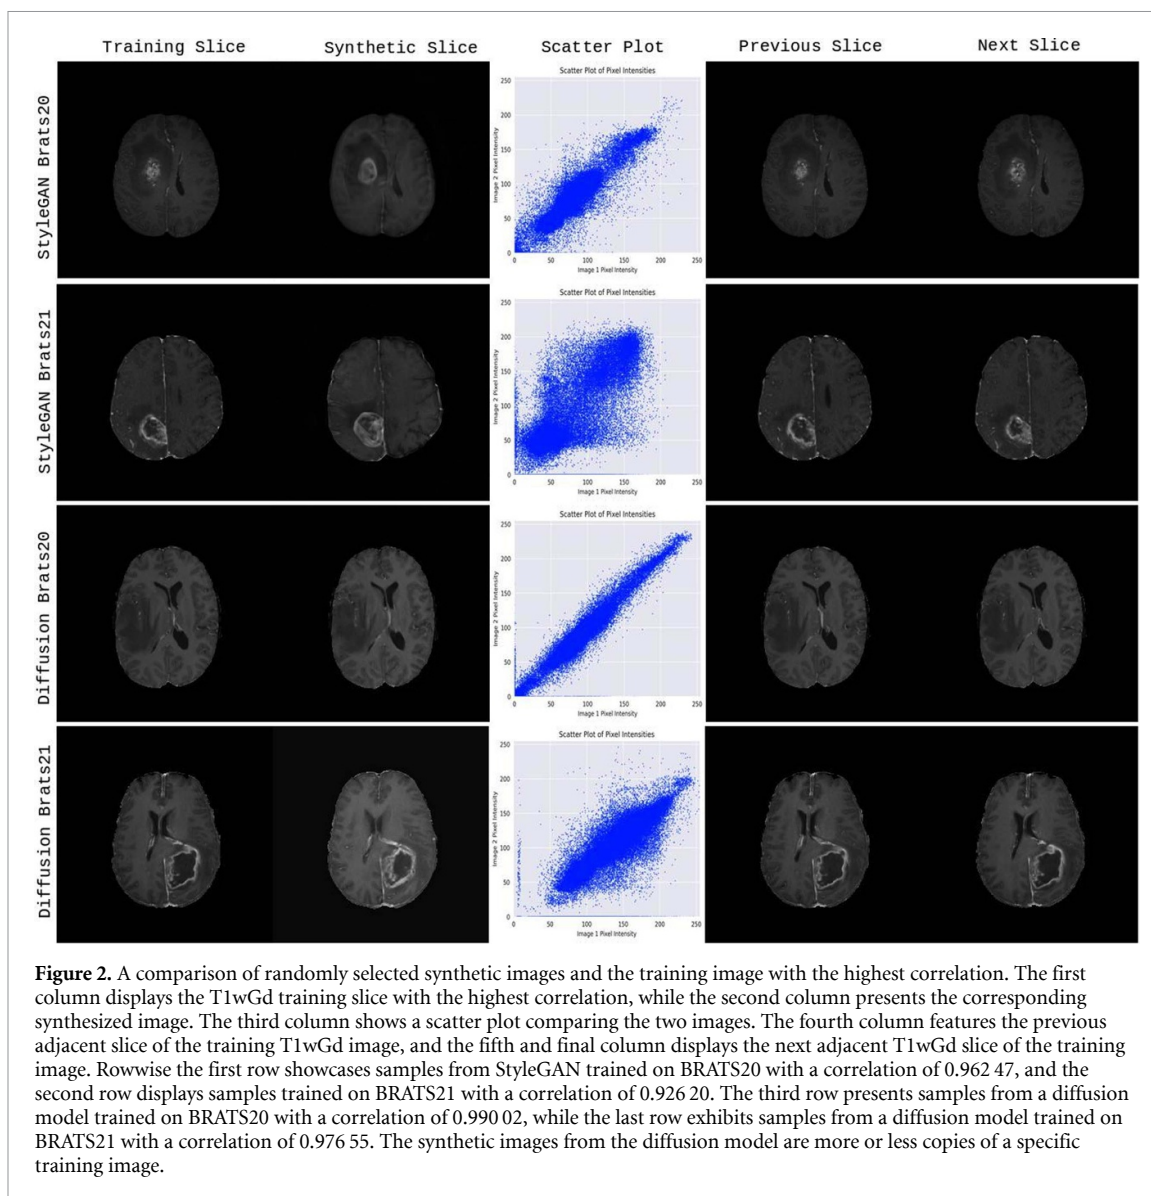
Another approach for quantifying memorization could be to use the feature vector in an autoencoder, which can work effectively even for rotated synthetic images for quantifying the memorization. An autoencoder maps images to a lower-dimensional feature space, typically between 100 to 1000 dimensions. While this approach is computationally efficient, the reduced dimensionality might not capture the fine details necessary for detecting memorization as effectively as pixel-wise correlation. Thus, the feature vector approach might not be as sensitive in identifying subtle memorization in generative models.

2.4. How do hyperparameters influence memorization?

In deep learning, the number of trainable parameters in a model is closely related to the likelihood of overfitting or memorization. To investigate how the size of generative models impacts memorization, smaller and larger versions of StyleGAN and the diffusion model were trained. For StyleGAN, size variations were achieved by modifying the capacity multiplier and the maximum number of feature maps. The default model used a capacity multiplier of 32 768 and a maximum of 512 feature maps. In the smaller model, these values were reduced to 16 768 and 256, respectively, while in the larger model, they were increased to 64 786 and 1024. For the diffusion model, size adjustments were made by varying the number of channels (filters), with the smaller model using 32 channels, the default model using 64, and the larger model using 96. Other hyperparameters for the diffusion model, including three residual blocks, 4000 diffusion steps, a linear noise schedule, learning rate, and batch size remained unchanged.

For CXR data, the default StyleGAN model had 57 335 748 trainable parameters (generator: 28 471 875, discriminator: 28 863 873), while the diffusion model had 31 050 885 trainable parameters. In the smaller StyleGAN model, the parameters were reduced to 47 542 337 (generator: 23 455 739, discriminator: 24 086 598), while the larger model had 76 583 792 parameters (generator: 40 260 785, discriminator: 36 323 007). Similarly, the smaller diffusion model had 7780 933 parameters, while the larger one had 69 809 861.

For brain MR data, which contained five channels, the default StyleGAN model had 59 962 308 trainable parameters (generator: 31 098 435, discriminator: 28 863 873), while the diffusion model had 41 123 331 parameters. The smaller StyleGAN model had 15 581 971 parameters (generator: 8343 757, discriminator: 7238 214), and the larger model had 220 400 166 parameters (generator: 105 533 217, discriminator: 114 868 977). The smaller diffusion model had 10 304 259 parameters, while the larger one had 92 457 219. The resulting synthetic images had random artefacts (not present for the default settings), and the training time was for this reason increased for training the smaller and larger models.



3. Results and analysis

In figures 1 and 4, we can see sample training images that were used to train the diffusion model and StyleGAN, along with randomly selected synthetic images from each model. Figure 2 compares synthetic images generated by StyleGAN and a diffusion model trained on BRATS20 and BRATS21 datasets while figure 5 compares images generated by the two generative models trained on CXR-5216 and CXR-1300 datasets. The comparison includes synthetic images, the training image with the highest correlation, a scatter plot, and adjacent training slices (for brain MRI). Notably, the diffusion model's images closely resemble specific training images, indicating higher data memorization compared to StyleGAN. Each row corresponds to different model-dataset combinations.

A total of 1000 synthetic images were randomly selected from each of our eight trainings (StyleGAN and diffusion model for BRATS20, BRATS21 and two CXR pneumonia datasets), and were used to calculate the correlation with all training images, and the highest correlations were saved. For BRATS each correlation was determined by combining all four modalities. To obtain a baseline, the correlations were also calculated between the training images and the test images from each dataset. Furthermore, the correlations were also calculated between the synthetic images and the test images. This was done to assess the level of correlation for data not used for training. From the plots in figures 3 and 6, it is evident that StyleGAN outperforms the diffusion model, as the latter produces images more similar to the training images. It is also clear that the degree of memorization is higher for the smaller BRATS20 dataset. Moreover from the CXR images, we can infer that the issue of memorization still persists, although not as severe. The plots demonstrate that correlation is an appropriate metric for investigating the memorization of training images.

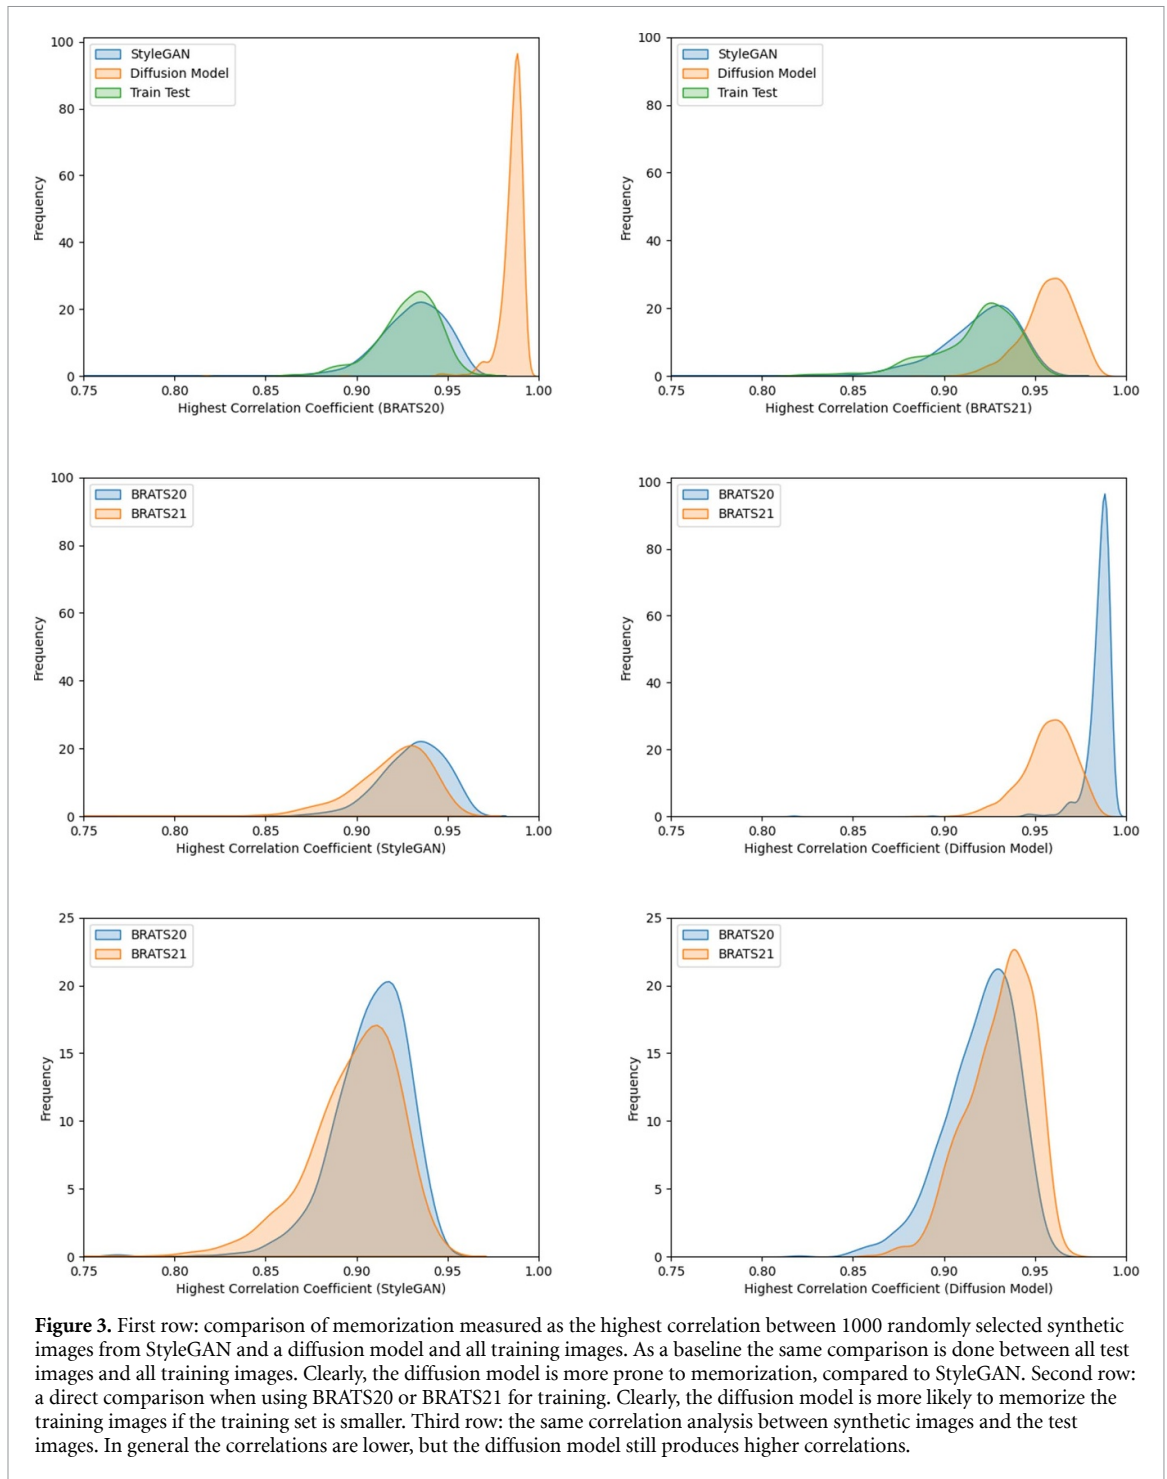



Figure 3. First row: comparison of memorization measured as the highest correlation between 1000 randomly selected synthetic images from StyleGAN and a diffusion model and all training images. As a baseline the same comparison is done between all test images and all training images. Clearly, the diffusion model is more prone to memorization, compared to StyleGAN. Second row: a direct comparison when using BRATS20 or BRATS21 for training. Clearly, the diffusion model is more likely to memorize the training images if the training set is smaller. Third row: the same correlation analysis between synthetic images and the test images. In general the correlations are lower, but the diffusion model still produces higher correlations.

Table 2 shows IS and FID for the different combinations of dataset and generative model. However, these metrics, including the commonly used FID, while indicative of the quality of generated images when lower, do not accurately represent the extent of data memorization relative to the training images, and we therefore also show the mean highest correlation. It is noted that generative models with better performance metrics, such as lower FID scores, often tend to memorize more data [13], highlighting a potential discrepancy between perceived image quality and actual model generalization [13].

Table 3 shows the mean highest correlation over 1000 synthetic images, for generative models with varying number of trainable parameters. The diffusion model tends to memorize more for larger models for x-ray images, while the other results are inconclusive. It should be mentioned that the synthetic images from the smaller and larger models are of lower quality compared to the medium (default) models, due to not performing an exhaustive hyperparameter search. This makes it difficult to draw any strong conclusions from this experiment.

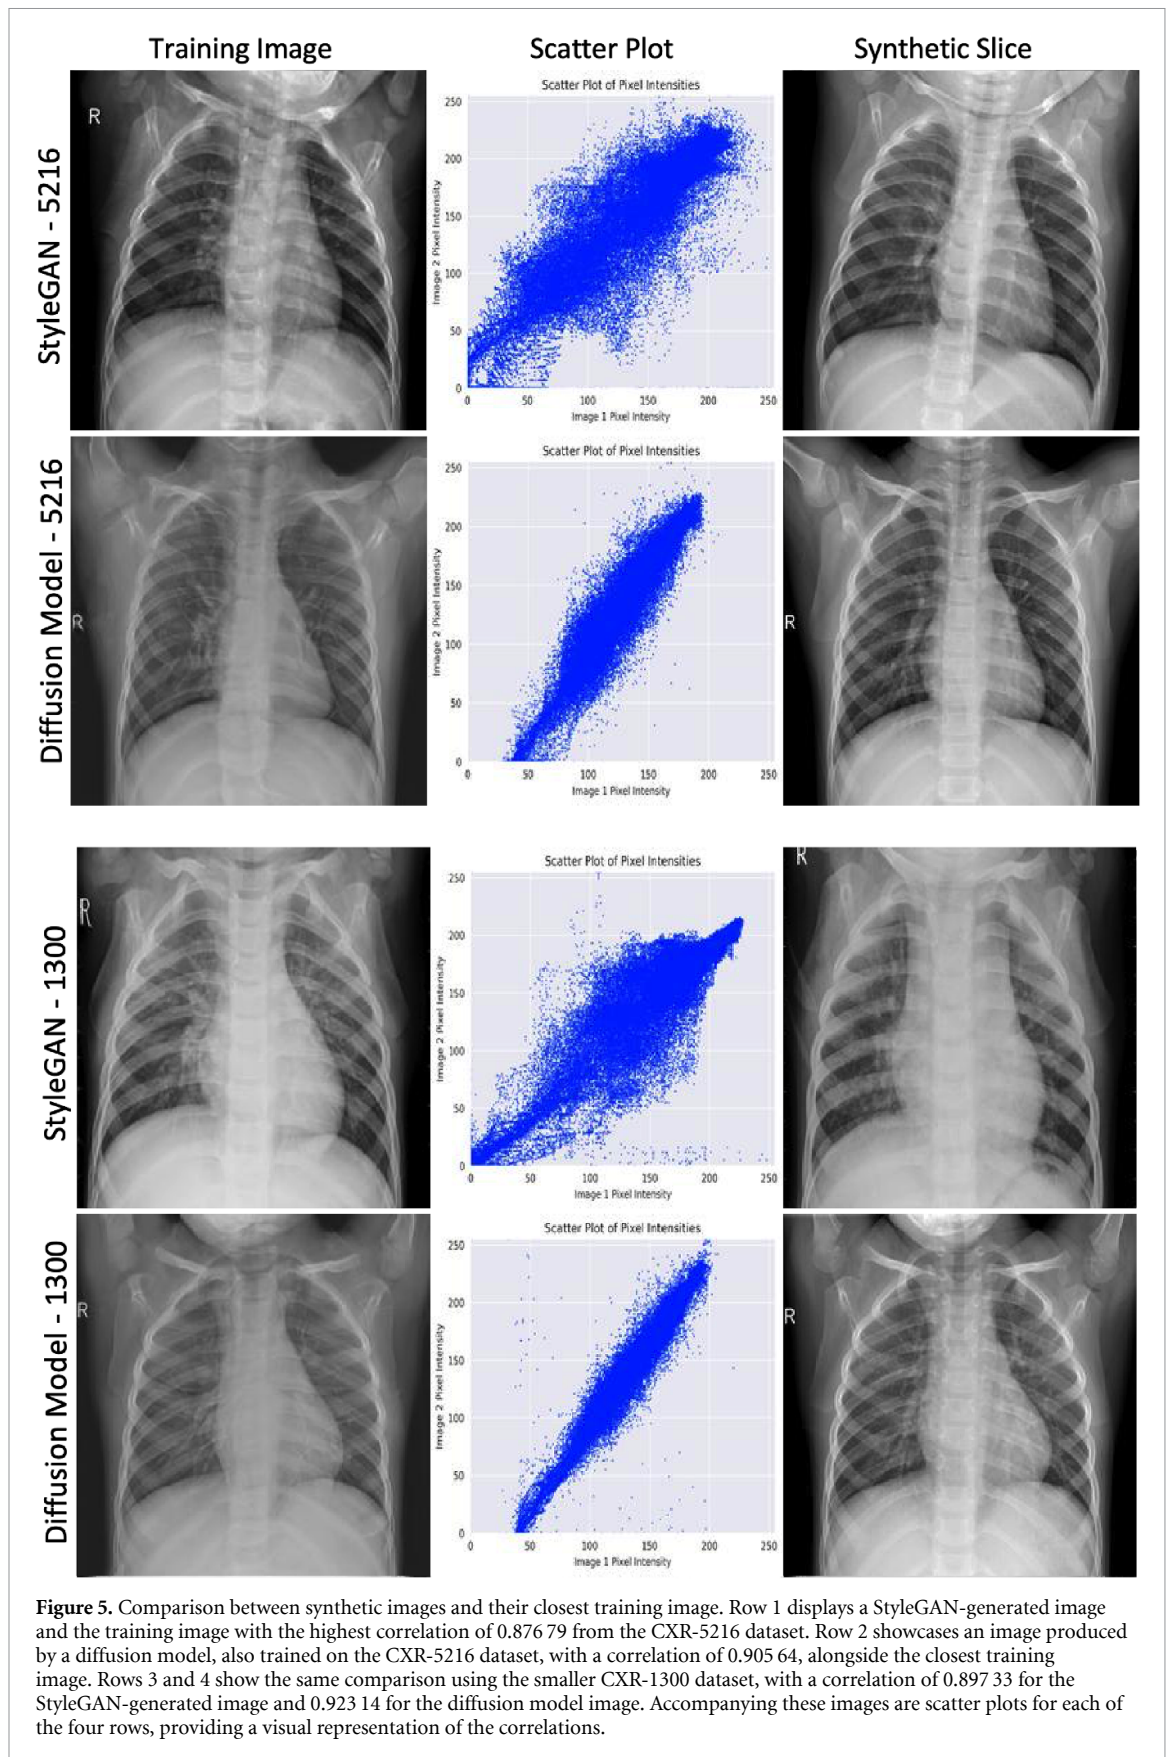


4. Discussion

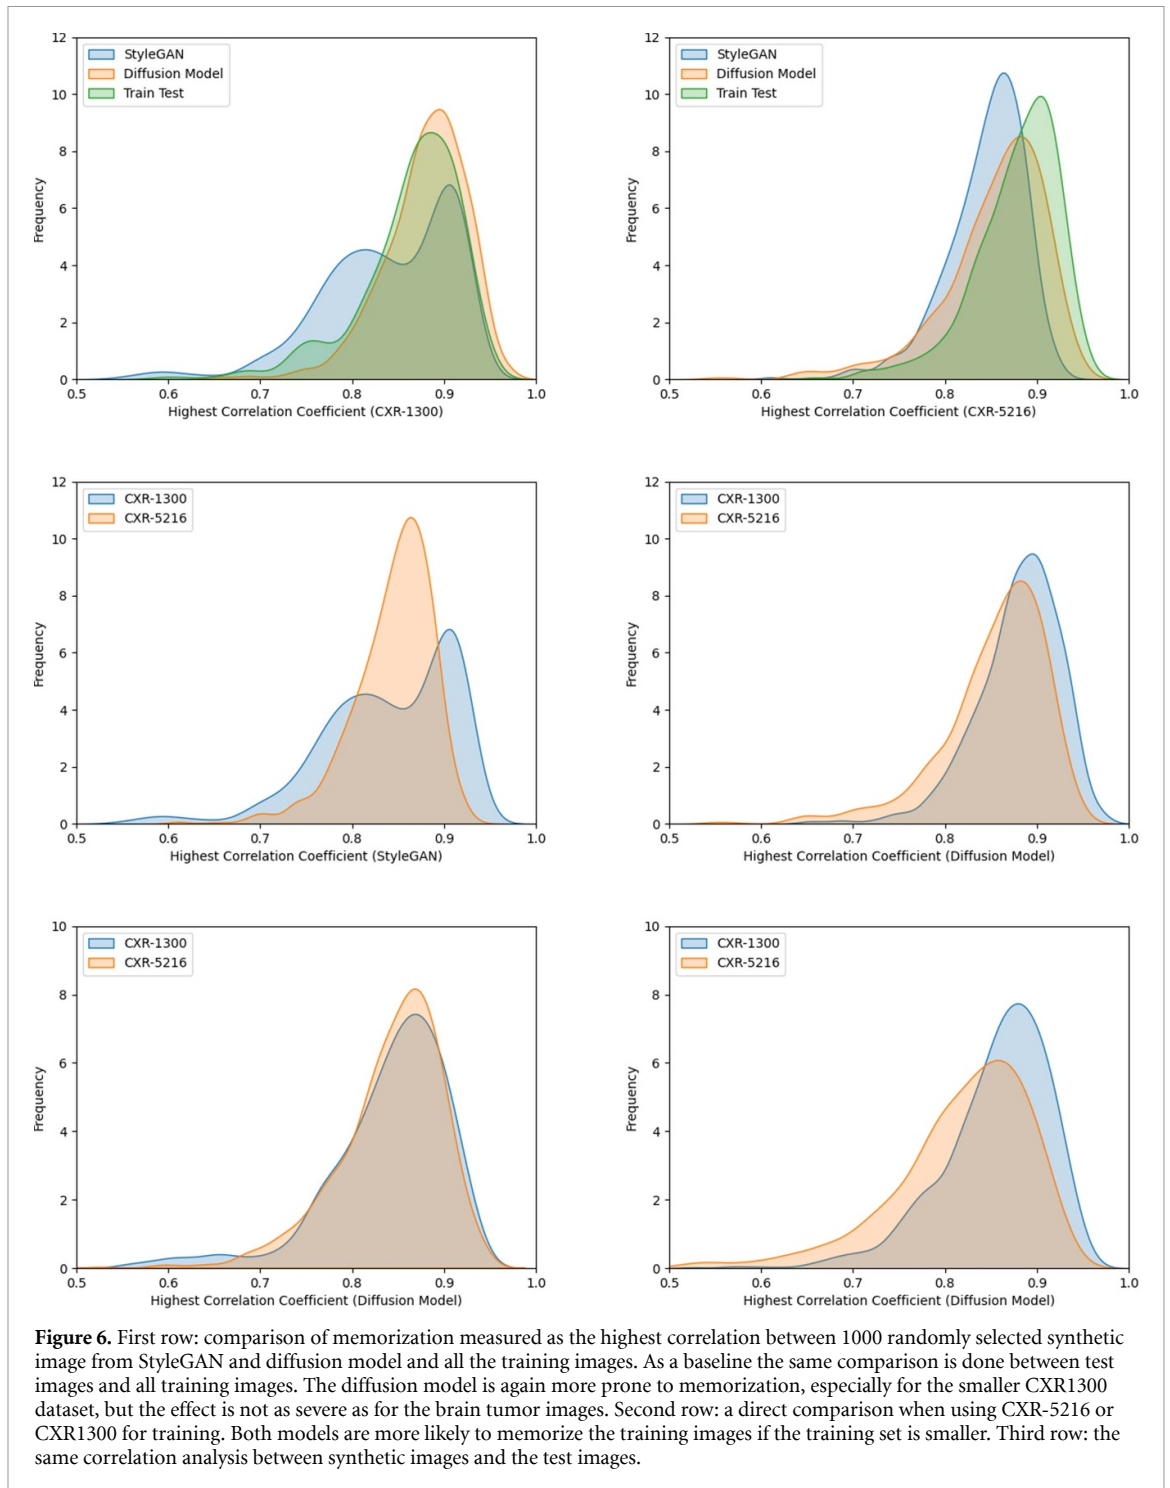
We have demonstrated that diffusion models can memorize the training images, corroborating previous results on this topic for non-medical images [13, 41] as well as medical images [15] (here no comparison with GANs was performed). The diffusion model tends to memorize (much) more than StyleGAN, even though the number of trainable parameters is lower (depending on if one includes the trainable parameters for the GAN discriminator or not). Memorization is problematic for medical imaging where privacy and GDPR is extremely important. Pinaya *et al* [33] have already shared synthetic brain volumes generated by their 3D diffusion model, and it should be investigated how similar these volumes are to the real brain volumes in UK biobank.

4.1. Effect of data size and data type

We show that the degree of memorization depends on the number of training images and type of data, as the problem is less severe for BRATS21 (1195 training subjects) compared to BRATS20 (313 training subjects). We see the same effect for the pneumonia data; a smaller training set leads to more severe memorization. The synthetic images produced from BRATS exhibit a much higher correlation compared to those produced from CXR pneumonia data. One possible explanation is that the BRATS generative models have more trainable parameters, as the MR images contain 5-channels compared to a single channel for CXR. Another explanation is that BRATS contains many slices in its training set that are highly correlated with each other,



as the dataset is originally 3D, while in the CXR pneumonia dataset the images are from unique subjects. Pan *et al* [31], Txurio *et al* [43] and Peng *et al* [32] also used 2D slices from 3D volumes, as 3D diffusion models are not yet very common. We expect that memorization is much more likely in these cases.



4.2. Effect of hyperparameters

One possible explanation for the observed memorization is the number of trainable parameters each model has (in relation to the number of training images). To study how the number of trainable parameters correlate with the degree of memorization, we trained smaller versions with fewer parameters and larger versions with more parameters, for both StyleGAN and the diffusion model. The correlation results for both models are inconclusive, but indicates that more trainable parameters can lead to more memorization for diffusion models. Training GANs and diffusion models is very time consuming due to the complexity of their architecture and the iterative nature of their training process. Given these constraints, conducting numerous experiments to fine-tune all hyperparameters becomes impractical, and we therefore only changed a few hyperparameters to make the models smaller or larger. This resulted in artefacts in the synthetic images, making it difficult to draw strong conclusions regarding how the number of trainable parameters affect memorization.

Table 2. A common comparison between GANs and diffusion models is to calculate Frechet inception distance and inception score for the synthetic images. However, these metrics do not consider memorization. Here we also show the mean highest correlation between 1000 synthetic images and all training images. Clearly the diffusion model memorizes more compared to StyleGAN.

Data	Model	FID	Inception score	Correlation
BRATS20	Diffusion model	15.854 17	2.351 91	0.985 10
	StyleGAN	44.856 24	2.341 84	0.930 37
BRATS21	Diffusion model	27.548 55	2.351 56	0.957 18
	StyleGAN	25.346 07	2.343 60	0.919 07
CXR-5216	Diffusion model	53.818 66	2.329 37	0.853 46
	StyleGAN	33.144 18	2.342 87	0.843 56
CXR-1300	Diffusion model	57.453 22	2.056 46	0.878 95
	StyleGAN	54.682 77	2.417 64	0.838 47

Table 3. A comparison of models trained with different number of trainable parameters (medium represents the default settings). The values represent the mean highest correlation for 1000 synthetic images. The mean correlation for the small BRATS20 diffusion model is very low. This is because the model was trained for the same number of epochs as the larger one, but the images produced are extremely poor.

Data	Model	Small	Medium	Large
CXR-1300	Diffusion model	0.863 63	0.878 95	0.885 19
	StyleGAN	0.852 99	0.838 47	0.843 36
BRATS20	Diffusion model	0.667 44*	0.985 10	0.977 31
	StyleGAN	0.947 56	0.930 37	0.926 39

4.3. Importance of detecting memorization

Medical images often contain identifiable information that can be linked to specific individuals, making handling and storage of this data a matter of significant concern. Mismanagement or unauthorized access to such data could lead to breaches of patient privacy, potentially resulting in legal consequences and reputational damage for medical institutions. Therefore, it is essential to implement strict security measures and adhere to privacy regulations when working with medical imaging data. In the majority of cases, medical data is deemed sensitive and is often restricted from being shared, even for research purposes. This is due to various concerns, including ethical considerations, privacy protection, and the sensitive nature of the information itself. Using synthetic images can resolve many issues and may be considered safe for sharing, as synthetic images do not belong to a specific person. However, a concern remains: what if the models generating these images are not creating unique content, but instead reproduce the training data?

For big generative models, one can wonder if they create genuinely new results or if they just copy and mix elements from their training data. By examining how these models remember, it can be understood how much they copy from the data they were trained on. Saharia *et al* [39] found that overfitting is not a problem in the model they trained model, and that more training could improve the performance of the diffusion model. This leads us to question if diffusion models perform better than previous methods because they remember more information.

4.4. Preventing memorization

There are several potential ways to prevent memorization when training generative models. For example, one may introduce a memorization penalty in the loss function [10]. Another approach is to only use augmented images (i.e. after applying rotation, flipping, elastic deformations, etc), and to never use the original images. However, this will make it harder and much more time consuming to evaluate how similar the synthetic images are to the original training images, if it is necessary to first apply image registration to be able to compare each image pair. Similar to Dar *et al* [15], one can train an autoencoder to compress each image or volume, and instead measure the similarity of the feature vectors. If the autoencoder is trained with augmentation, it should be able to detect a high similarity even if two images have different orientation. As previously mentioned, the number of elements in the feature vector may however be too low to detect memorization of fine details. As proposed in our related work [2], to fine-tune generative models pre-trained on large open datasets can potentially lead to less memorization, compared to training these models from scratch. Another approach for preventing memorization is to use techniques similar to differential privacy (DP) [1], but recent work demonstrated that diffusion models trained with DP did not converge [13]. To memorize the training images is an overfitting problem, and common regularization techniques such as dropout and weight decay can potentially help. Recent work has demonstrated that model gradient similarity [42] can be very helpful to automatically penalize memorization.

4.5. Implications for practical clinical applications

Our findings highlight the need for careful guidelines to prevent generative models from memorizing training data, especially in clinical settings where privacy and data security are critical. If a generative model remembers too much, it may accidentally recreate identifiable patient information, posing privacy risks. For this reason, researchers should only share the synthetic images and not the generative models. It is rather easy to investigate if a synthetic image is similar to a training image, but it is substantially harder to guarantee that a generative model does not contain sensitive data (even if DP or other privacy-protecting techniques have been used). By following these practices, researchers and medical institutions can use synthetic data confidently in research and diagnostics. This makes data more accessible, improves model reliability, and ensures that privacy rules are respected. In this way, synthetic data can be shared and used more widely in clinical settings, supporting innovation while keeping patient information safe.

5. Conclusion

This study highlights an important issue with using diffusion models (and to some extent GANs) for medical images: the problem of these models remembering and possibly copying the data they were trained on. This is especially true for small datasets or datasets with very similar images. This not only causes problems with how well the models work, but also raises privacy concerns because there's a chance these models might recreate images that can be linked to real sensitive data. We have here studied memorization for brain and chest images, but the problem is equally important for all types of medical images.

We found that memorization is more noticeable in certain datasets, like BRATS, than in others like the CXR pneumonia dataset which has more varied images. This tells us that how a model behaves depends a lot on the kind of data it learns from and how much data there is. Considering these points, it is really important for people who make and use these models, especially in medical imaging, to find the right balance. They need to use the models' power for good while making sure they handle private medical information correctly. Looking ahead, we should keep looking for ways to prevent these models from memorizing too much and think carefully about the ethical side of using them in healthcare.

Data availability statement

The synthetic brain MR images used in this study are openly available at the following URL/DOI: <https://datahub.aida.scilifelab.se/10.23698/aida/synthetic/brgandi> [2].

Acknowledgment

Training StyleGAN, and sampling from the diffusion models, was performed using the supercomputing resource Berzelius (752 Nvidia A100 GPUs) provided by the National Supercomputer Centre at Linköping University, Sweden. It was donated by the Knut and Alice Wallenberg Foundation.

This research was supported by the ITEA/VINNOVA project ASSIST (Automation, Surgery Support and Intuitive 3D visualization to optimize workflow in IGT Systems, 2021-01954), LiU Cancer and the Åke Wiberg foundation. Anders Eklund has previously received graphics hardware from Nvidia.

ORCID iDs

Muhammad Usman Akbar  <https://orcid.org/0000-0002-3248-5132>

Anders Eklund  <https://orcid.org/0000-0001-7061-7995>

References

- [1] Abadi M, Chu A, Goodfellow I, McMahan H B, Mironov I, Talwar K and Li Z 2016 Deep learning with differential privacy *Proc. 2016 ACM SIGSAC Conf. on Computer and Communications Security* pp 308–18
- [2] Akbar M U, Larsson M, Blystad I and Eklund A 2024 Brain tumor segmentation using synthetic MR images-A comparison of GANs and diffusion models *Sci. Data* **259**
- [3] Ali H, Murad S and Shah Z 2023 Spot the fake lungs: generating synthetic medical images using neural diffusion models *Artificial Intelligence and Cognitive Science: 30th Irish Conf., AICS 2022, Munster, Ireland, December 8–9, 2022, Revised Selected Papers* pp 32–39
- [4] Bakas S, Akbari H, Sotiras A, Bilello M, Rozycki M, Kirby J S, Freymann J B, Farahani K and Davatzikos C 2017 Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features *Sci. Data* **4** 170117

- [5] Bakas S, Reyes M, Jakab A, Bauer S, Rempfler M, Crimi A, Shinohara R T, Berger C, Ha S M and Rozycki M 2018 Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the BRATS challenge (arXiv:1811.02629)
- [6] Bakas S, Akbari H, Sotiras A, Bilello M, Rozycki M, Kirby J, Freymann J, Farahani K and Davatzikos C 2017 Segmentation labels and radiomic features for the pre-operative scans of the TCGA-GBM collection *The Cancer Imaging Archive* (<https://doi.org/10.7937/K9/TCIA.2017.KLXWJJ1Q>)
- [7] Bakas S, Akbari H, Sotiras A, Bilello M, Rozycki M, Kirby J, Freymann J, Farahani K and Davatzikos C 2017 Segmentation labels and radiomic features for the pre-operative scans of the TCGA-LGG collection *The Cancer Imaging Archive* (<https://doi.org/10.7937/K9/TCIA.2017.GJQ7R0EF>)
- [8] Beaulieu-Jones B K, Wu Z S, Williams C, Lee R, Bhavnani S P, Byrd J B and Greene C S 2019 Privacy-preserving generative deep neural networks support clinical data sharing *Circ. Cardiovascular Qual. Outcomes* **12** e005122
- [9] Borji A 2019 Pros and cons of GAN evaluation measures *Comput. Vis. Image Underst.* **179** 41–65
- [10] Borji A 2022 Pros and cons of GAN evaluation measures: new developments *Comput. Vis. Image Underst.* **215** 103329
- [11] Brock A, Donahue J and Simonyan K 2018 Large scale GAN training for high fidelity natural image synthesis (arXiv:1809.11096)
- [12] Çalli E, Sogancioglu E, Ginneken B, Leeuwen K G and Murphy K 2021 Deep learning for chest x-ray analysis: a survey *Med. Image Anal.* **72** 102125
- [13] Carlini N, Hayes J, Nasr M, Jagielski M, Sehwal V, Tramer F, Balle B, Ippolito D and Wallace E 2023 Extracting training data from diffusion models (arXiv:2301.13188)
- [14] Chen C, Wang H, Chen Y, Yin Z, Yang X, Ning H, Zhang Q, Li W, Xiao R and Zhao J 2023 Understanding the brain with attention: a survey of transformers in brain sciences *Brain-X* **1** e29
- [15] Dar S U H, Ghanaat A, Kahmann J, Ayx I, Papavassiliou T, Schoenberg S O and Engelhardt S 2023 Investigating data memorization in 3D latent diffusion models for medical image synthesis (arXiv:2307.01114)
- [16] Dhariwal P and Nichol A 2021 Diffusion models beat GANs on image synthesis *Advances in Neural Information Processing Systems* vol 34 pp 8780–94
- [17] Eilertsen G, Tsirikoglou A, Lundström C and Unger J 2021 Ensembles of GANs for synthetic training data generation *ICLR 2021 Workshop on Synthetic Data Generation* **10.48550/arXiv.2104.11797**
- [18] El Emam K, Mosquera L, Jonker E and Sood H 2021 Evaluating the utility of synthetic COVID-19 case data *JAMIA Open* **4** o0ab012
- [19] Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A and Bengio Y 2014 Generative adversarial networks *Advances in Neural Information Processing Systems (NeurIPS)*
- [20] Heusel M, Ramsauer H, Unterthiner T, Nessler B and Hochreiter S 2017 GANs trained by a two time-scale update rule converge to a local Nash equilibrium *Advances in Neural Information Processing Systems* vol 30
- [21] Ho J, Jain A and Abbeel P 2020 Denoising diffusion probabilistic models *Advances in Neural Information Processing Systems* vol 33 pp 6840–51
- [22] Karras T, Laine S and Aila T 2019 A style-based generator architecture for generative adversarial networks *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition* pp 4401–10
- [23] Karras T et al 2021 Alias-free generative adversarial network *Advances in Neural Information Processing Systems (NeurIPS)*
- [24] Kazerouni A, Aghdam E K, Heidari M, Azad R, Fayyaz M, Hacıhaliloglu I and Merhof D 2023 Diffusion models in medical imaging: a comprehensive survey *Med. Image Anal.* **textbf88** 102846
- [25] Kermany D S, Goldbaum M, Cai W, Valentim C C S, Liang H, Baxter S L, McKeown A, Yang G, Wu X and Yan F 2018 Identifying medical diagnoses and treatable diseases by image-based deep learning *Cell* **172** 1122–31
- [26] Khader F et al 2023 Denoising diffusion probabilistic models for 3D medical image generation *Sci. Rep.* **13** 7303
- [27] Larsson M, Akbar M U and Eklund A 2022 Does an ensemble of GANs lead to better performance when training segmentation networks with synthetic images? (arXiv:2211.04086)
- [28] Menze B H, Jakab A, Bauer S, Kalpathy-Cramer J, Farahani K, Kirby J, Burren Y, Porz N, Slotboom J and Wiest R 2014 The multimodal brain tumor image segmentation benchmark (BRATS) *IEEE Trans. Med. Imaging* **34** 1993–2024
- [29] Moghadam P A, Van Dalen S, Martin K C, Lennerz J, Yip S, Farahani H and Bashashati A 2023 A morphology focused diffusion probabilistic model for synthesis of histopathology images *Proc. IEEE/CVF Winter Conf. on Applications of Computer Vision* pp 2000–9
- [30] Packhäuser K, Folle L, Thamm F and Maier A 2022 Generation of anonymous chest radiographs using latent diffusion models for training thoracic abnormality classification systems (arXiv:2211.01323)
- [31] Pan S, Wang T, Qiu R L J, Axente M, Chang C -W, Peng J, Patel A B, Shelton J, Patel S A and Roper J 2023 2D medical image synthesis using transformer-based denoising diffusion probabilistic model *Phys. Med. Biol.* **68** 105004
- [32] Peng J, Chen G, Saruta K and Terata Y 2023 2D brain MRI image synthesis based on lightweight denoising diffusion probabilistic model *Med. Imaging Process Technol.* **6** 2518
- [33] Pinaya W H L, Tudosiu P-D, Dafflon J, Da Costa P F, Fernandez V, Nachev P, Ourselin S and Cardoso M J 2022 Brain imaging generation with latent diffusion models *Deep Generative Models: Second MICCAI Workshop, DGM4MICCAI 2022, Held in Conjunction With MICCAI 2022, Singapore, September 22, 2022, Proc.* pp 117–26
- [34] Rajotte J-F, Bergen R, Buckeridge D L, El Emam K, Ng R and Strome E 2022 Synthetic data as an enabler for machine learning applications in medicine *Iscience* **25** 105331
- [35] Ramesh A, Dhariwal P, Nichol A, Chu C and Chen M 2022 Hierarchical text-conditional image generation with clip latents (arXiv:2204.06125)
- [36] Ramesh A, Pavlov M, Goh G, Gray S, Voss C, Radford A, Chen M and Sutskever I 2021 Zero-shot text-to-image generation *Int. Conf. on Machine Learning* pp 8821–31
- [37] Rankin D, Black M, Bond R, Wallace J, Mulvenna M and Epelde G 2020 Reliability of supervised machine learning using synthetic data in health care: model to preserve privacy for data sharing *JMIR Med. Inform.* **8** e18910
- [38] Rombach R, Blattmann A, Lorenz D, Esser P and Ommer B 2022 High-resolution image synthesis with latent diffusion models *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition* pp 10684–95
- [39] Saharia C, Chan W, Saxena S, Li L, Whang J, Denton E L, Ghasemipour K, Gontijo Lopes R, Karagol Ayan B and Salimans T 2022 Photorealistic text-to-image diffusion models with deep language understanding *Advances in Neural Information Processing Systems* vol 35 pp 36479–94

- [40] Salimans T, Goodfellow I, Zaremba W, Cheung V, Radford A and Chen X 2016 Improved techniques for training GANs *Advances in Neural Information Processing Systems (NeurIPS)* vol 29
- [41] Somepalli G, Singla V, Goldblum M, Geiping J and Goldstein T 2023 Diffusion art or digital forgery? Investigating data replication in diffusion models *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition* pp 6048–58
- [42] Szolnoky V, Andersson V, Kulcsar B and Jörnsten R 2022 On the interpretability of regularisation for neural networks through model gradient similarity *Advances in Neural Information Processing Systems (NeurIPS)*
- [43] Txurio M S, Román K L-L, Marcos-Carrión A, Castellote-Huguet P, Santabárbara-Gómez J M, Oliver I M'ia and Ballester M A G 2023 Diffusion models for realistic CT image generation *Int. KES Conf. on Innovation in Medicine and Healthcare* pp 335–44
- [44] Yi X, Walia E and Babyn P 2019 Generative adversarial network in medical imaging: a review *Medical Image Analysis* **58** 101552