Contents lists available at ScienceDirect



Computer Methods and Programs in Biomedicine

journal homepage: www.elsevier.com/locate/cmpb

Inflating 2D convolution weights for efficient generation of 3D medical images



Yanbin Liu^{a,1}, Girish Dwivedi^{b,c}, Farid Boussaid^d, Frank Sanfilippo^e, Makoto Yamada^f, Mohammed Bennamoun^{g,*}

^a School of Computing, Australian National University, Canberra, ACT, AU

^b Harry Perkins Institute of Medical Research, The University of Western Australia, Perth, WA, AU

^c Cardiology Department, Fiona Stanley Hospital, Perth, WA, AU

^d Department of Electrical, Electronic and Computer Engineering, The University of Western Australia, Perth, WA, AU

^e School of Population and Global Health, The University of Western Australia, Perth, WA, AU

^fOkinawa Institute of Science and Technology, Okinawa, JP

^g Department of Computer Science and Software Engineering, The University of Western Australia, Perth, WA, AU

ARTICLE INFO

Article history: Received 12 April 2023 Revised 13 June 2023 Accepted 19 June 2023

Keywords: Weight inflation Generative adversarial networks Efficient neural networks Three-dimensional medical images

ABSTRACT

Background and Objective: The generation of three-dimensional (3D) medical images has great application potential since it takes into account the 3D anatomical structure. Two problems prevent effective training of a 3D medical generative model: (1) 3D medical images are expensive to acquire and annotate, resulting in an insufficient number of training images, and (2) a large number of parameters are involved in 3D convolution. *Methods*: We propose a novel GAN model called *3D Split&Shuffle-GAN*. To address the 3D data scarcity issue, we first pre-train a two-dimensional (2D) GAN model using abundant image slices and inflate the 2D convolution weights to improve the initialization of the 3D GAN. Novel 3D network architectures are proposed for both the generator and discriminator of the GAN model to significantly reduce the number of parameters while maintaining the quality of image generation. Several weight inflation strategies and parameter-efficient 3D architectures are investigated. *Results*: Experiments on both heart (Stanford AIMI Coronary Calcium) and brain (Alzheimer's Disease Neuroimaging Initiative) datasets show that our method leads to improved 3D image generation quality (14.7 improvements on Frchet inception distance) with significantly fewer parameters (only 48.5% of the baseline method). *Conclusions*: We built a parameter-efficient 3D medical image generation model. Due to the efficiency and effectiveness, it has the potential to generate high-quality 3D brain and heart images for real use cases.

© 2023 Elsevier B.V. All rights reserved.

1. Introduction

With the availability of large-scale annotated datasets like ImageNet [1], convolution neural networks (CNNs) have achieved unprecedented success in computer vision [2]. Benefiting from CNNs, medical imaging research has made great advancements in the classification [3], segmentation [4–6], detection [7], reconstruction [8], and registration [9] of two-dimensional (2D) medical images. However, 3D medical image research lags behind due to the lack of large-scale 3D medical image datasets. As a result of the complex collection procedure, expert annotation, privacy concerns and patient consent, it is challenging to build a large-scale, 3D medical dataset similar to ImageNet.

* Corresponding author.

E-mail address: mohammed.bennamoun@uwa.edu.au (M. Bennamoun).

¹ Most work was done at Harry Perkins Institute of Medical Research.

One widely-used solution for the data deficit of medical images is Generative Adversarial Networks (GANs) [10]. These networks create high-quality synthetic images to mimic realistic data distributions. An example is using GANs with Wasserstein distance and perceptual loss for low-dose computed tomography (CT) image denoising [11]. Perceptual loss cannot be directly used for 3D medical images due to the lack of interpretable pre-trained 3D models. A cyclic loss GAN was used by [12] to reconstruct MRI images. Using cycle-consistent GANs, [13] translated magnetic resonance (MR) images to CT images. Albeit effective in mitigating the datadeficit challenge, most existing GANs-based methods are designed for 2D medical image generation. Therefore, they do not incorporate information about the 3D anatomical structure [14]. Various medical applications require the 3D anatomical structure, including calcium scoring [15,16] of cardiac CT Coronary Angiograms (CTCAs), and brain tumor segmentation [17,18]. Unfortunately, there are two practical issues that hinder the effective training of the 3D medical generative model, preventing the use of GANs in 3D medical imaging.

First of all, there are usually insufficient 3D medical images to train effective 3D generative models. The effective training of 3D CNNs with natural videos relies on large-scale datasets, such as Moments in Time [19] with 1 million short videos, and Kinetics [20] with 750k video clips. In comparison, medical datasets contain far fewer 3D images. For example, Stanford AIMI Coronary Calcium (COCA) dataset [21] only contains 787 CTCAs. To generate 3D images, [22] used 991 brain MRI images from the Alzheimers Disease Neuroimaging Initiative (ADNI) dataset [23]. Training an effective 3D generative model is difficult with such small datasets of 3D medical images.

Secondly, 3D convolution layers have a large number of parameters, making training time slow and prone to overfitting. As a result of 3D convolution, weight parameters take on an additional dimension. For example, the conventional 3×3 convolution expands to $3 \times 3 \times 3$ in the 3D case. Adding the third dimension allows the modeling of the 3D anatomical structure, but it also involves the introduction of an excessive number of parameters and computations, resulting in slower training. Moreover, the model is prone to overfitting due to the contrast between the large number of parameters and the small number of 3D training images.

To address the above two problems, we propose a novel GAN model, dubbed **3D Split&Shuffle-GAN** for effective and efficient 3D medical image generation. The proposed model improves existing state-of-the-art GANs (e.g., StyleGAN2 [24]) from two perspectives: training strategy and network architecture (see Fig. 1).

The proposed training strategy takes advantage of the availability of 2D image slices to train a 2D GAN model. It then inflates the 2D weights to initialize the 3D GAN model. As the 3D GAN model is initialized with informative 2D weights, it can focus more on 3D anatomy, which results in a better generation of 3D images. By design, the 2D GAN shares a similar architecture as the 3D GAN, with the exception of the additional convolution dimension (e.g., 3×3 convolution vs. $3 \times 3 \times 3$ convolution). This enables the 2D weights to be seamlessly expanded to 3D using the weight inflation technique [25]. Since the original inflation was designed for classification models rather than generative models, we evaluate five new inflation variants through extensive experiments to determine the most suitable one for the task of 3D image generation.

For the network architecture, we devise novel Channel Split&Shuffle modules to improve both the generator and discriminator networks. For the generator, since the state-of-the-art style-based models (e.g., StyleGAN2) incorporate style vectors into convolution weights as a modulated convolution, efficient convolution operations (e.g., depthwise separable convolution [26] or group convolution [2]) cannot be directly adopted. This is mitigated by our Split&Shuffle module, which splits the feature channels into



Fig. 1. The general pipeline of our 3D generative model, which includes our contribution to both the training strategy (inflate 2D weights, colored in blue) as well as the network architecture (Split&Shuffle GAN, colored in orange). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

two equal branches and performs a modulated 3D convolution for each branch. Then, the output channels are concatenated and shuffled to encourage feature exchanges. With this design, the number of parameters of the generator is reduced by a factor of 2. For the discriminator, the number of parameters is further reduced by nearly a factor of 4 by replacing one of the $3 \times 3 \times 3$ convolutions with $1 \times 1 \times 1$ convolution. Although the number of parameters for both the generator and discriminator is significantly reduced, the devised modules achieve a much better performance than the original one. Under the extremely data-deficit challenges of generating 3D medical images, our parameter-efficient model is less likely to overfit than the original model.

To demonstrate the effectiveness of the training strategy and network architecture, we investigated five novel weight inflation variants as well as five network design choices on the heart dataset (COCA). In addition, we performed experiments on the brain dataset (ADNI) to demonstrate the general applicability of our method.

To summarize, this paper makes the following contributions to 3D medical image generation:

- A novel 3D Split&Shuffle-GAN model for 3D medical image generation is proposed, and new inflation strategies are developed to facilitate training of 3D medical generation models.
- Parameter-efficient Channel Split&Shuffle modules are developed for both the generator and discriminator networks, which reduces the number of parameters (by a factor of at least 2) and improves generation quality (FID).
- We conducted comprehensive experiments to verify the effectiveness of the inflation strategy and network architecture. We achieved state-of-the-art performance on both the heart and brain datasets.

2. Related work

2.1. Generative models for medical imaging

The most popular model for generating synthetic images is the generative adversarial networks (GANs) [10]. The GANs model synthesizes realistic images from a random noise variable and uses a discriminator to distinguish between the synthesized images and the realistic images. The distribution of the synthesized images gradually approaches the distribution of real images with the alternating training of the generator and discriminator. State-of-the-art GANs use the style-based generation technique [24,27], in which style vectors are generated (for controlling the style of image generation) from a mapping network.

Providing annotations to large numbers of images in the field of medical imaging is a challenging task. The use of GANs is thus naturally adopted to solve a number of medical problems [3,11-13,28-33], such as classification, segmentation, registration, low dose CT denoising, and MR to positron emission tomography (PET) synthesis. GANs were used by [3] to generate synthetic CT images for data augmentation to enhance liver lesion classification performance. RefineGAN [12] proposes a cyclic consistency loss for the modified variant of the deeper generator and discriminator networks to deal with the compressed sensing magnetic resonance imaging (CS-MRI) reconstruction problem. To improve the conventional GANs for the low dose CT (LDCT) denoising task, [11] employed two practical methods, namely Wasserstein distance and perceptual loss. A-CycleGAN [13] makes use of variational autoencoding (VAE), attention, and cycle-consistent generative adversarial network (CycleGAN) to improve existing MR-to-CT image translation algorithms. [31] proposed an effective adversarial U-Net architecture along with different normalization techniques to solve the MRI to PET image synthesis task.

Despite the wide range of models and GAN variants proposed for medical imaging problems, most of them only focus on generating 2D images, disregarding the 3D anatomical structure. Only a few attempts have been made to generate 3D images. Leveraging an α -GAN, [22] utilizes the variational autoencoder (VAE) and GAN to generate 3D synthetic brain MRI images. [4] proposed a segmentation-guided style-based generative adversarial network (SGSGAN) for synthesizing full-dose PET images, where a stylebased generator is directly used for style modulation. [34] proposed a hierarchical shape-perception network (HSPN) for 3D brain reconstruction (point cloud) from a single incomplete image. In contrast, our method generates 3D medical images with only random variables as input. By extending StyleGAN2's 2D convolutions to 3D convolutions, [35] used 3D-StyleGAN to generate 3D brain MRI images. A comprehensive review of the usage of GANs in 3D data can be found in [14]. Since most existing methods lift 2D GANs models to 3D in a straightforward manner, the number of parameters increases significantly, making it challenging to train the model effectively. In this paper, we propose both effective training strategies and efficient model architectures to generate 3D medical images using 3D GANs.

2.2. Training 3D convolution neural networks

A multitude of research effort has been directed toward 3D CNNs in the field of natural images, especially for the spatiotemporal analysis of videos. The main idea is to introduce a third convolution dimension ($k \times k \times k$) to capture the temporal dependencies for video applications such as action recognition [25]. The training of 3D CNN models usually relies on large-scale video datasets, e.g., Kinetics-700 [20] with 750k video clips, Moments in Time [19] with 1 million short videos. However, due to the high annotation costs, patient consent issues, and expert annotation challenges, creating 3D medical image datasets of similar scale is not feasible. As a result, training 3D medical models is challenging.

Using degenerated 2D spatial information, another line of work contributes to initializing 3D convolution weights by utilizing beneficial priors. For example, [25] proposed an inflation strategy to stack 2D weights for the 3D weights initialization. The video vision Transformer is trained using the central frame initialization strategy in [36]. To our knowledge, no similar initialization technique has been explored in 3D medical GANs. On the one hand, the third dimension in video analysis corresponds to temporally varying frames, while the third dimension in medical images describes the 3D anatomical structure. On the other hand, the interplay between the discriminator and generator makes the training process more complex than that of classification models. In this paper, we consider both the 3D anatomical structure and the interplay between the discriminator and the generator to facilitate the 3D GAN training and architecture design.

2.3. Parameter-efficient 3D convolution neural networks

3D convolution neural networks are challenged by the large number of parameters included by the additional third dimension. There are two main approaches to addressing this issue: tensor decomposition and efficient module design. In tensor decomposition, the low-rank tensor decomposition algorithms are applied to re-calculate the convolution weights, thereby compressing the network and reducing the number of parameters. For example, Tensor Train has been used in [37], CANDECOMP/PARAFAC (CP) decomposition is applied in [38,39], and Tucker decomposition is adopted in [40]. In spite of their mathematical soundness, these methods require specific re-implementation of existing convolution operations and cannot take advantage of the latest hardware acceleration (e.g., the NVIDIA cudnn library). In efficient module design, various parameter-efficient modules (e.g., bottleneck [41], group convolution, depthwise separable convolution, and pointwise convolution) are devised to replace the original module. These efficient modules are re-arranged and combined to form different network architectures. In MobileNet [26], for example, depthwise separable convolutions are used to construct a lightweight deep architecture for mobile devices. SqueezeNet [42] combines pointwise convolution and regular convolution to form a Fire block. The computation cost of ShuffleNet [43] is reduced by using pointwise group convolution. A comprehensive analysis of these modules can be found in [44].

All the above parameter-efficient designs are based on classification models, which cannot be directly and easily adopted in 3D generative models, such as StyleGAN2. StyleGAN2's style modulation mechanism will be destroyed if these modules are trivially adopted. To address this issue, we propose customized 3D modules for the style-based generative models to enable parameter-efficient generation of 3D medical images.

3. Methods

3.1. Preliminary of 3D medical image generation

3.1.1. Overview of stylegan2 model

Mapping Network A key difference between style-based generative models (e.g., StyleGAN2) and previous GANs is the introduction of the mapping network f. Specifically, given a latent code $z \in \mathbb{Z}, f : \mathbb{Z} \to W$ first produces a vector $w \in W$. The learned affine transform A is then applied to w to obtain the generator's per-layer style vectors s.

Generator In the generator, original StyleGAN [27] directly utilizes the style vectors for adaptive instance normalization (AdaIN) on the feature maps, which will cause characteristic artifacts such as droplets. To mitigate these unrealistic artifacts, StyleGAN2 incorporates the style vectors into the weight modulation (Mod) operation, then applies the demodulation (Demod) to serve as the instance normalization.

Modulation:
$$w_{ijk} = s_i \cdot w_{ijk}$$
 (1)

Demodulation:
$$w''_{ijk} = w'_{ijk} / \sqrt{\sum_{i,k} w'_{i,j,k}^2 + \epsilon}$$
 (2)

where w, w', w'' are the original, modulated and demodulated convolution weights, s_i is the style vector corresponding to the *i*-th feature map, *j*, *k* iterate the output feature maps and the spatial resolution, ϵ is a small constant to avoid numerical issues. In the above modulation and demodulation operations, the style vectors are directly entangled with convolution weights, which removes the characteristic artifacts while retaining the style controllability. However, this also impedes the straightforward modification to the convolution layers, e.g., depthwise separable convolution (details in Section 3.3).

Discriminator The discriminator of StyleGAN2 introduces a minibatch standard deviation layer to calculate the deviation of a minibatch and concatenates it to the original feature maps. This reduces the dependency on a minibatch to encourage diverse generations.

3.1.2. 3D Medical image generation

The StyleGAN2 was originally designed for 2D natural image generation. To apply it to the generation of 3D medical images, a straightforward approach [35] is by lifting all the 2D convolution operations to the 3D convolution operations, e.g., expanding the 3×3 convolutions to the $3 \times 3 \times 3$ convolutions. Albeit simple,



Fig. 2. The overall architecture of the proposed 3D Split&Shuffle-GAN. It is composed of a Mapping Network, a Generator, and a Discriminator. Mapping Network maps a latent variable *z* to the style vector space W and produces the per-layer style vectors with a learned affine transform (A). The generator starts from a constant 3D input. Then it controls the 3D generation styles with the per-layer style vectors and adds details from the per-channel scaled (B) noise input. Weight Mod&Demod incorporates the style vector into the convolution operation $(w_3^{(1)}, w_3^{(2)}, w_3^{(3)}, \ldots$ represent the 3D convolution weights for first, second, third,...layers). Discriminator tries to differentiate the real 3D images from the generated fake 3D images. Inside both the Generator and Discriminator, we devise novel **Channel Split & Shuffle** modules for parameter-efficient 3D convolution operations, which are customized for the style-based generation framework.

this approach will significantly increase the number of parameters, thereby posing two practical issues: (1) it will require a large number of 3D images for training; otherwise, the model suffers overfitting and mode collapse issues², and (2) the largely increased parameter number will slow down training and generation. However, no existing methods have simultaneously addressed both issues. Therefore, it is non-trivial to improve StyleGAN2 for 3D medical image generation.

In this paper, we deal with the issues from the **training strategy** (Weight Inflation in Section 3.2) and **network architecture** (Split&Shuffle in Section 3.3) perspectives and propose an efficient 3D generative model as shown in Fig. 2.

3.2. Inflating 2D convolution weights

This section discusses how to design a training strategy to generate synthetic 3D medical images and overcome the issue of data scarcity. Transfer learning [46] is a widely-used technique to overcome the data paucity of the target task by employing additional external datasets. As an example, ImageNet [1] is used as an external dataset for a variety of tasks, including object detection, and segmentation. In MinGAN [47], knowledge is transferred from GANs to domains with few images. It is currently difficult to apply the transfer learning technique for 3D GAN models in medical research due to the lack of large datasets and effective transfer learning strategies. 3DSeg-8 dataset [48] has been aggregated from eight datasets to facilitate transfer learning between 3D medical images for liver segmentation and nodule classification. However, there are only tens to hundreds of 3D images of organs/tissues in each dataset in 3DSeg-8, which makes GANs in-

capable of transferring detailed knowledge. Moreover, the interplay between the generator and discriminator poses a different transfer learning challenge compared to classification tasks with a single network.

Although direct transferring from external datasets is not feasible for 3D medical image generation, we note an interesting and helpful observation: **2D slices in a medical dataset are several magnitudes larger than 3D images**. As an example, COCA [21] contains only 787 CTCA images, but these CTCA images contain 39,281 2D slices. Hence, the number of these 2D slices is sufficient to train a 2D generative model such as StyleGAN2. Since 2D and 3D generative models have distinct weights, it is not possible to transfer weights directly from 2D to 3D generative models. As a result, we are naturally drawn to another technique called weight inflation [25], which enables effective 3D network training from the pre-trained 2D weights.

Weight inflation was first introduced in [25] for the design and training of 3D video action recognition networks. The method has been applied to both CNNs and Transformers-based 3D models [25,36,49]. In technical terms, it extends/inflates/copies the 2D convolution weights along the third dimension (e.g., temporal dimension in the video) to provide a more favorable initialization for the 3D convolution networks. To ensure feasible inflation from 2D weights to 3D models, the 2D and 3D networks must have the same basic structure except for the additional third dimension (e.g., the 3×3 and $3 \times 3 \times 3$ convolutions should have the same number of channels). **To our knowledge, this technique has not been applied to medical image analysis for effective 3D medical image generation**.

Considering the application differences between video action recognition and 3D medical image generation, we propose five customized inflation strategies to facilitate the training of 3D Style-GAN2. Here, we set the size of convolution weights to be $3 \times 3 \times 3$, but our strategies can be easily adapted to other weight sizes.

² Overfitting means that on small datasets, the discriminator overfits to the training examples and training starts to diverge [45]. Mode collapse means that the generator only produces a limited set of outputs.

Let $w_2 \in \mathbb{R}^{C_l \times C_0 \times 3 \times 3}$ denote the pre-trained 2D convolution weights and $w_3 \in \mathbb{R}^{C_l \times C_0 \times 3 \times 3 \times 3}$ denote the corresponding 3D convolution weights (C_l represents the number of input channels, C_0 represents the number of output channels). At first, we initialize w_3 from a random Gaussian $\mathcal{N}(0, 0.1)$. Then, the weight w_3 is modified by the following inflation strategies:

• **Inflate-1**: only inflating 1 center dimension.

$$w_3[:,:,1,:,:] = w_2. \tag{3}$$

• Inflate-2: inflating 2 dimensions.

 $w_3[:,:,0,:,:] = w_2, \quad w_3[:,:,1,:,:] = w_2.$ (4)

• Inflate-3: inflating all 3 dimensions.

$$w_{3}[:,:,0,:,:] = w_{2}, \quad w_{3}[:,:,1,:,:] = w_{2}, w_{3}[:,:,2,:,:] = w_{2}.$$
(5)

• **Inflate-ASC**: inflating the axial, sagittal, and coronal planes of 3D medical views.

$$w_{3}[:,:,1,:] = w_{2}, \quad w_{3}[:,:,:,1,:] = w_{2}, w_{3}[:,:,:,1] = w_{2}.$$
(6)

• **Inflate-NWI**: the negative weight initialization (NWI) is used here, with the center dimension owing a larger value and the other dimensions owing negative values (T = 3).

$$w_{3}[:,:,i,:] = \alpha_{i} * w_{2}, \ \alpha_{i} = \begin{cases} \frac{2T-1}{T}, \text{ if } i = 1\\ -\frac{1}{T}, \text{ otherwise}. \end{cases}$$
(7)

By design, different inflation strategies offer diverse ways of reusing the 2D weights: the reusing degree increases from Inflate-1 to Inflate-3; Inflate-ASC considers the anatomical views; Inflate-NWI modifies Inflate-1 with more attention on the center dimension and negative weights on others. Intuitively, these inflation strategies introduce helpful 2D structure priors through weight initialization, which significantly reduces the training burden of the 3D convolution weights. In this way, by focusing more on the third dimension (for anatomy learning), the generative model is able to generate high-quality 3D images quickly and efficiently.

3.3. Efficient 3D architecture design

Although inflation strategy mitigates the lack of 3D data in training the 3D GAN model, it still suffers from the large number of model parameters. In this section, we address this issue from the perspective of efficient 3D architecture design.

We observe that most of the parameters in the 3D neural network architecture originate from the 3D convolution operation, which extends the 2D convolution weights to 3D, to model the 3D contextual and anatomical structures (e.g., lifting 3×3 weights to $3 \times 3 \times 3$). Existing efficient 3D architecture designs mainly focus on parameter-efficient 3D convolution. On the one hand, factorized high-order CNNs are proposed with different tensor decomposition algorithms such as Tensor-Train (TT) [37], CP decomposition [38,39], Tucker Decomposition [40]. These methods compress networks and reduce their parameters by applying low-rank tensor decompositions to high-order weights. On the other hand, driven by the requirements of mobile devices, various parameter-efficient convolution variants have been devised and combined into efficient architectures, such as group convolution [2], bottleneck [50], depthwise separable convolution [26]. Group convolution divides the channels into groups and performs convolution only within each group. Bottleneck was introduced in ResNet [50] to reduce the number of channels of the 3×3 convolution by wrapping it with two 1×1 convolutions. With depthwise separable convolutions, the standard convolutions are factorized into a depthwise convolution (i.e., a group convolution with a group number equal to a channel number) followed by a 1×1 pointwise convolution.

The above designs were developed to improve the efficiency of various applications, e.g., HO-CPConv [39] for spatiotemporal facial emotion analysis, MobileNet [26] for image classification and object detection, and 3D-MobileNet [44] for video action recognition. Despite this, the efficient 3D GANs architecture is rarely studied, especially for the state-of-the-art StyleGAN2 model. We attribute this to two possible reasons: (1) The StyleGAN2 architecture is more delicate than classification models, hindering the straightforward adoption of existing modules such as tensor decomposition, group convolution, or depthwise separable convolution. Specifically, the style vectors are absorbed in the modulation and demodulation operations (Equations 1,2), which sets up hurdles for existing modules. (2) As a result of the interplay between the discriminator and generator, GANs training is difficult. It is non-trivial to directly use the same modules for the discriminator and generator to achieve the best performance.

Based on the above analysis, we propose a unique design customized to the StyleGAN2 model for parameter-efficient generation (Fig. 3). For the generator, since the direct adoption of existing efficient modules (e.g., group convolution, depthwise separable convolution) will break the entangled structure of the convolution weights and style vectors (Equations 1,2), we equally Split the feature maps, using the channel split operation, to create two branches. The modulation and demodulation operations for the style vectors and convolution weights are individually applied to each branch. Afterwards, the outputs of the two branches are concatenated, and the Channel Shuffle operation is performed, allowing information to be shared between two channels. If the Channel Shuffle operation is not performed, the generator is considered to be two independent networks. Channel Shuffle enables hybrid and diverse pattern combinations across branches to facilitate image generation quality. For the discriminator, as neither modulation nor demodulation is applied, there is more flexibility in improving the design. Therefore, we devised two asymmetric branches with $3 \times 3 \times 3$ and $1 \times 1 \times 1$ convolutions. The $1 \times 1 \times 1$ convolution leads to a further parameter reduction while compromising the local spatial structure. But this is rectified by the Channel Shuffle operation, which exchanges information by shuffling the feature maps of $1 \times 1 \times 1$ and $3 \times 3 \times 3$ convolutions.

Considering C = 32 in Fig. 3, both the input and output feature maps are of size $32 \times H \times W \times D$. Without splitting, the size of the convolution weight is $32 \times 32 \times 3 \times 3 \times 3$. With splitting, the input map is split into two $16 \times H \times W \times D$ branches, each undergoing a 3D convolution (weight size $16 \times 16 \times 3 \times 3 \times 3$). The outputs of two branches are concatenated to get the feature map of size $32 \times H \times W \times D$. So, the total size of the convolution weight is $2 \times 16 \times 16 \times 3 \times 3 \times 3$, which means the generator enjoys a parameter reduction of 2. Similarly, the discriminator has a total weight size of $16 \times 16 \times 3 \times 3 \times 3 + 16 \times 16 \times 1 \times 1 \times 1$, which means it enjoys a parameter reduction of nearly 4 (3.857).

We also propose several other parameter-efficient convolution architectures as baselines to verify the Split&Shuffle design's effectiveness. As a result of the modulation and demodulation constraint in the generator, several baselines only modify the discriminator (i.e., D only). All the model variants are listed below:

• **Group Convolution (D only)**, which replaces the convolution in the discriminator with a group convolution.



Fig. 3. The proposed Channel Split&Shuffle Convolution Modules for the Generator and Discriminator. The Discriminator module has two differences from the Generator module: the style Mod&Demod and $1 \times 1 \times 1$ Convolution. Overall, the proposed module reduces the number of parameters by a factor of 2 in the Generator and nearly by a factor of 4 in the Discriminator. Here, w_{11} and w_{12} denote the 3D convolution weights of the left and right branches, which are used to perform the Mod&Demod operation in Equations 1–2.

- **Depthwise Separable Convolution (D only)**, which replaces the convolution in the discriminator with a depthwise separable convolution.
- **Split&Shuffle Convolution (D only)**, which applies the Split&Shuffle module only on the discriminator.
- **Split Convolution**, which applies the channel split without channel shuffle.
- Split&Shuffle Convolution, which is our final design.

4. Results

4.1. Datasets and evaluation metrics

4.1.1. Datasets

Stanford AIMI Coronary Calcium (COCA) [21] dataset is used for heart CTCA images. COCA contains 787 3D coronary CT images. Each 3D image has a different number (ranging from 27 to 156) of 2D slices on the axial plane, and they add up to 39,281 axial slices in total. For brain MRI images, we used the Alzheimers Disease Neuroimaging Initiative (ADNI) dataset [23]. Specifically, we used 991 T1 structural images from the Cognitively Normal (CN) research group. MR images from non-brain areas were removed by the dataset provider using the software FreeSurfer's³ recon-all function. The processed MR images have 256 slices from all three planes.

4.1.2. Evaluation metric

In order to assess the quality of generated images, GANs usually use the FrEchet inception distance (FID) [51] metric, which compares the feature distributions of real and generated images. By default, the Inception V3 network [52] pre-trained with 2D natural images was deployed for feature extraction. However, the method cannot be directly applied to 3D images. As such, taking into account the 3D medical structure, we measured the FID scores on the center slices of axial, sagittal and coronal planes, i.e., *FID-ax*, *FID-sag*, and *FID-cor*. Lastly, we averaged the three FID scores to obtain *FID-avg* as an overall measurement.

However, the FID alone cannot evaluate 3D medical image generation comprehensively. The reasons are twofold: (1) the inception model only accepts 2D image slices of a specific plan, and (2) the inception model is pre-trained on natural images with a large gap with medical images. Therefore, we also adopted the widelyused metrics MS-SSIM, PSNR, and the t-Distributed Stochastic Neighbour Embedding (t-SNE) to evaluate the performance.

³ https://surfer.nmr.mgh.harvard.edu/fswiki

4.2. Implementation details

4.2.1. Pre-processing

Both the COCA and ADNI images were resized to 64×64 and 128×128 using bilinear interpolation in the sagittal and coronal planes. We aligned the axial slice number of the COCA dataset to 32 either by consecutive slice sampling or by zero padding. For the ADNI dataset, we resized the axial slice number to 64 directly. As a result, the image resolution of COCA and ADNI datasets are $32 \times 64 \times 64$ ($32 \times 128 \times 128$) and $64 \times 64 \times 64$, respectively. The Hounsfield Unit values were clipped to [-250, 650] and then normalized to [-1, 1].

4.2.2. Stylegan2 architecture

The base layer (Const in Fig. 2) for the COCA dataset was set to $2 \times 4 \times 4$, followed by 5 upsampling stages. For the ADNI dataset, the base layer was set to $4 \times 4 \times 4$, followed by 5 upsampling stages. The total number of convolution channels was set to 32, with each branch owing 16. The feature dimension of the mapping network was set to 64.

4.2.3. Hyper-parameters

We followed the StyleGAN2 configuration for training, with the following exceptions: $\gamma = 0.0512$ for R_1 regularization, minibatch = 128, learning rate lr = 0.0025 for training from scratch, lr = 0.002 for inflation initialization. For 2D GANs pre-training, the models are fed with 25,000K images in total. However, the 3D GANs were trained with 5,000K images since 3D models are slow to train.

4.3. Inflation and architecture analysis

4.3.1. Inflation

The first step is to conduct comprehensive experiments to verify that inflation strategies are effective for initializing the 3D generative model. A 2D StyleGAN2 model is pre-trained using all the 39,281 axial slices to obtain the 2-dimensional convolution weights. Since the number of images is sufficient, the 2D model achieves an FID of 7.71. This means that the pre-trained 2D weights capture rich slice-level contextual information to generate highquality 2D slices. Thus, it verifies the rationale behind inflating 2D pre-trained weights for 3D generative models.

Starting from the same pre-trained 2D weights, we apply all the proposed inflation strategies as well as a "No Inflate" baseline to initialize and train the 3D StyleGAN2 model. The results are shown in Table 1. It can be seen that three inflation strategies outperform the "No Inflate" baseline (the other two are comparable),



Fig. 4. Performance of different inflation variants during the training process (kimgs). Among all, "Inflate-1" generally obtains good FID scores on all three views. Red arrows show the best FID iterations (kimgs) for "No Inflate" and "Inflate-1".

 Table 1

 Performance of different inflation variants (lower numbers are better).

Model	FID-ax	FID-sag	FID-cor	FID-avg
No Inflate	104	92	112	102.7
Inflate-1	60	83	76	73.0
Inflate-2	68	133	110	103.7
Inflate-3	89	119	110	106.0
Inflate-ASC	79	108	83	90.0
Inflate-NWI	85	87	83	85.0

indicating that inflation strategies are generally effective as favorable initialization methods for 3D generative models. "Inflate-1", which only initializes one center dimension, achieves the best performance (FID-avg) among all inflation variants. Performance gradually degrades as more weights are initialized from 2D weights ("Inflate-2" and "Inflate-3"). We hypothesise that the overly inflated 3D weights will prevent the model from freely learning the third-dimensional anatomical structure. "Inflate-ASC" and "Inflate-NWI" perform slightly worse than "Inflate-1".

To understand the training dynamics of various inflation strategies, we plot the FID scores with respect to the training iteration (measured by the number of 1,000 training images, i.e., kimgs) in Fig. 4. In general, "Inflate-1" achieves the best performance during training, consistent with Table 1. Most inflation variants achieve significantly better FID scores than the baseline at the start of training (i.e., kimgs=0), demonstrating the effectiveness of 2D weight priors in training StyleGAN2. "Inflate-NWI" initially performs worse because it modifies the original 2D weights, but in the end, it outperforms the baseline since the prior informative 2D weights have made a big difference.

In Fig. 5, we randomly generated the axial/sagittal/coronal slices of CT images for both the "No Inflate" baseline and our best variant "Inflate-1" to intuitively investigate why the inflation strategy works. Specifically, we show the generated images from the initial training iteration (i.e., the dashed lines in Fig. 4) and the best training iteration (i.e., the red arrows in Fig. 4). It is easy to observe that with inflation as a favorable initialization, the generated images already show meaningful anatomical structures even before training (e.g., FID-ax=201). In contrast, the randomly initialized "No Inflate" generates blurry meaningless images before training (FID-ax=440). This comparison provides an intuitive explanation for the working mechanism of the inflation strategy: with effective inflation, the 3D generative model can inherit meaningful 2D anatomical priors for better subsequent training. Furthermore, starting from better initial weights, the inflated model ("Inflate-1") is trained to achieve superior generative performance compared to the "No Inflate" baseline (e.g., FID-ax=60 vs FID-ax=104).



Fig. 5. Generated images from the initial training (dashed lines in Fig. 4) and the best (red arrows in Fig. 4) training iterations. (a) With random weights as an initialization (No Inflate), the initially generated images are meaningless. (b) With inflated weights as a favorable initialization (Inflate-1), the initial generated images already show basic anatomical structures. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

StyleGAN2's discriminator and generator have different structures by design. This motivates us to examine how inflation affects the discriminator and the generator. Specifically, we selected the three best inflation variants ("Inflate-1", "Inflate-ASC", and "Inflate-NWI") and performed three sets of experiments (Table 2): inflating the generator only (G), inflating the discriminator only (D), and inflating both the generator and discriminator (G&D, default). This analysis reveals two observations: (1) inflating the entire model (G&D) always achieves the best performance, (2) the generator plays a more important role in the inflation strategy, which is rea-

Table 2

Inflation on Generator (G) and Discriminator (D).

Model	Module	FID-ax	FID-sag	FID-cor	FID-avg
Inflate-1	G	105	82	88	91.7
	D	85	94	104	94.3
	G&D	60	83	76	73.0
Inflate-ASC	G	93	96	94	94.3
	D	123	113	127	121.0
	G&D	79	108	83	90.0
Inflate-NWI	G	85	91	94	90.0
	D	90	92	98	93.3
	G&D	85	87	83	85.0

Table 3

Parameter and performance comparison on different architectures. Models with "-D" mean only the discriminator owns the modified architecture.

Model	#param	FID-ax	FID-sag	FID-cor	FID-avg
Baseline	0.600M	104	92	112	102.7
Group-D	0.434M	80	130	105	105.0
Depthwise-D	0.367M	67	136	117	106.7
Split&Shuffle-D	0.416M	84	121	118	107.7
Split	0.291M	109	151	146	135.3
Split&Shuffle	0.291M	85	113	91	96.3

sonable because the generator is responsible for generating images with the style vectors to control the generation.

4.3.2. Architecture

Our next step was to conduct experiments to comprehensively investigate efficient 3D architectural designs. At first, all model variants, including the baseline, are trained from scratch without weight inflation. The FID values and the number of parameters are shown in Table 3. The models with a "-D" suffix apply efficient modules only on the discriminator, thus only leading to a slightly reduced parameter number (e.g., 0.434M vs. 0.600M). "Group-D", "Depthwise-D", and "Split&Shuffle-D" have a different number of parameters because the architectures are different. By contrast, the proposed "Split&Shuffle" design reduces more than half of the parameters compared to the baseline (0.291M vs 0.600M).

With regards to generation performance, the proposed "Split&Shuffle" architecture achieved the better FID-avg with the least number of parameters, proving its efficiency and effectiveness. All the "-D" models produce similar results to the baseline, indicating that modifying just the discriminator has little influence on the generation quality. Note that owing the same number of parameters, "Split" performs much worse than the "Split&Shuffle" design. As a result, the Channel Shuffle in Fig. 3 plays a crucial role in ensuring performance.

Then, we examined how efficient architectures can be combined with inflation strategies for further performance enhancement. Specifically, we adopted the best inflation strategy "Inflate-1" and applied it to all the models in Table 3. The results are shown in Table 4. In this experiment, each individual model has its own pre-trained 2D weights due to the differences in architecture. Table 4 demonstrates that all models achieved good FID-2D values, once again verifying the rationale for inflating informative 2D weights to train 3D GANs. Since the number of 2D image slices is sufficient for 2D pre-training, "Baseline" with the largest number of parameters achieved the best FID-2D. However, for the final 3D training, our "Split&Shuffle" model achieves the best performance (FID-avg) with the least number of parameters. Compared with the "Baseline" model, "Split&Shuffle" reduces the FID-avg by 14.7 with only 48.5% of the parameters. As for the discriminator-only variants with the suffix "-D", they were much worse than the "Inflate-1" baseline. Without channel shuffle operation, "Split" achieved the Computer Methods and Programs in Biomedicine 240 (2023) 107685



Fig. 6. Generated images of higher resolution (128×128) from the Baseline and our Split&Shuffle method. Baseline sometimes shows anatomically inconsistent regions (marked in red boxes). Our method can generate images with better anatomic structure and possible calcium slices (marked in blue boxes), which will be helpful for downstream tasks. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

worst performance, again showing the indispensable role of channel shuffle in our architecture design.

4.3.3. Image resolution

To show that our method can generate high-resolution images for practical application, we increased the resolution of COCA to 128×128 . We set the base layer to $1 \times 4 \times 4$, followed by six upsampling stages. The model capacity was also increased by using 64 convolution channels.

The quantitative results are shown in Table 5. Compared with the Baseline (same as Table 3), Split&Shuffle achieves much better performance with fewer model parameters. The visualization of the generated image slices is shown in Fig. 6. The generated image slices by our method show a more feasible heart anatomy structure and higher image quality. In addition, since COCA contains coronary calcium, our method generated image slices with possible calcium, which are more realistic.

4.4. Comparison with state-of-the-art methods

Finally, we compared the performance of our method with the published 3D generative models in Table 6 (COCA) and Table 7 (ADNI). The comparison methods included the following 3D generation baselines:

- 3D-WGAN-GP [53], a 3D extension of Wasserstein GAN with Gradient Penalty to alleviate training instability.
- 3D-VAE-GAN [54], consisting of an encoder, a decoder and a discriminator.
- 3D- α -GAN [55], applying the code discriminator and encoder on top of the conventional GANs to alleviate the collapse and blurriness.
- 3D- α -WGAN-GP [22], utilizing both the α -GAN structure and the WGAN-GP loss.
- 3D-StyleGAN2 [35], a direct 3D extension of the 2D StyleGAN2 model.

On both the heart and brain datasets, the proposed method outperforms all the state-of-the-art methods by a large margin, demonstrating its effectiveness. Among all comparison methods, the first four baselines have a much greater number of parameters than our method but achieved inferior performance. Although 3D-StyleGAN2 has approximately twice as many parameters as our method, it still performed much worse. Because the 3D medical image generation lacks sufficient 3D training images, most baselines are short of sufficient 3D training images, leading to inferior

Table 4

Model performance when inflation strategy ("Inflate-1") is applied to different architectures. FID_2D denotes the pre-trained 2D model performance.

Model	#param	FID-2D	FID-ax	FID-sag	FID-cor	FID-avg
Baseline	0.600M	7.7	60	83	76	73.0
Group-D	0.434M	10.9	78	76	76	76.7
Depthwise-D	0.367M	14.0	87	129	84	100.0
Split&Shuffle-D	0.416M	10.6	95	112	101	102.7
Split	0.291M	12.2	120	144	106	123.3
Split&Shuffle	0.291M	11.6	46	66	63	58.3

Table 5

Model performance comparison of higher resolution (128×128) images on COCA dataset. Baseline is the same as Table 3.

Model	#param	FID-ax	FID-sag	FID-cor	FID-avg
Baseline	2.628M	92	103	81	92.0
Split&Shuffle	1.414M	55	67	34	52.0

Table 6

Comparison with state-of-the-art methods on COCA (heart) dataset.

Model	#param	FID-ax	FID-sag	FID-cor	FID-avg
3D-WGAN-GP	1.399M	191	223	191	201.7
3D-VAE-GAN	6.389M	196	281	272	249.7
3D-α-GAN	2.811M	91	88	80	86.3
3D-α-WGAN-GP	2.811M	70	86	69	75.0
3D-StyleGAN2	0.600M	104	92	112	102.7
Split&Shuffle (Ours)	0.291M	46	66	63	58.3

Table 7

Comparison with state-of-the-art methods on ADNI (brain) dataset.

FID-ax	FID-sag	FID-cor	FID-avg
161	161	231	184.3
167	122	224	171.0
73	77	99	83.0
74	72	87	77.7
114	89	99	100.7
65	67	82	71.3
	FID-ax 161 167 73 74 114 65	FID-ax FID-sag 161 161 167 122 73 77 74 72 114 89 65 67	FID-ax FID-sag FID-cor 161 161 231 167 122 224 73 77 99 74 72 87 114 89 99 65 67 82

performance. As an alternative, our method uses both the weight inflation and Split&Shuffle designs to mitigate the reliance on a large number of 3D images for training.

Generation quality and generation diversity Except for FID scores, we considered two widely-used evaluation metrics: PSNR and MS-SSIM. Specifically, PSNR is calculated between the real and generated images to evaluate the generation quality. Following [22], MS-SSIM is calculated on pairs of generated images to evaluate the generation diversity (a smaller value means better diversity). The results are shown in Table 8. Our method achieved the best PSNR and MS-SSIM performance on both the heart and brain datasets, demonstrating that our method can generate high-

Table 8

PSNR and MS-SSIM evaluations on ADNI (brain) and COCA (heart) datasets. PSNR was calculated between real and generated images to evaluate generation quality (higher values are better). MS-SSIM was calculated on pairs of generated images to evaluate generation diversity (lower values are better).

	COCA (heart)		ADNI (brain	n)
Model	PSNR ↑	MS-SSIM ↓	PSMR ↑	MS-SSIM ↓
3D-WGAN-GP	16.22	0.9988	27.11	0.9945
3D-VAE-GAN	16.38	0.8436	27.03	0.9719
3D-α-GAN	16.19	0.8288	27.42	0.8750
$3D-\alpha$ -WGAN-GP	15.82	0.7916	27.71	0.8678
3D-StyleGAN2	16.04	0.8362	27.80	0.8999
Split&Shuffle (Ours)	16.56	0.7513	28.57	0.8488



Fig. 7. t-Distributed Stochastic Neighbour Embedding on COCA dataset.

quality 3D medical images with better diversity. We also note that the PSNR scores on the heart dataset are smaller than that on the brain dataset. This is due to the large variations among the unaligned heart images. In contrast, the brain images are aligned and exhibit smaller variations.

t-Distributed Stochastic Neighbour Embedding (t-SNE) To better understand the distributions of generated and real images, we performed t-SNE on real images, our method, $3D-\alpha$ -WGAN-GP, and 3D-StyleGAN2. The visualization of the COCA dataset is shown in Fig. 7. Although the distributions of both our method and $3D-\alpha$ -WGAN-GP approach the real images, our method is closer to the real images. The distribution of 3D-StyleGAN2 is far from the real images, consistent with its large FID score. As to the visualization



Fig. 8. t-Distributed Stochastic Neighbour Embedding on ADNI dataset.



Fig. 9. Generated images from the initial training and the best training on the brain ADNI dataset. Our method shows good initial brain anatomy, especially for the axial view.

of the ADNI dataset in Fig. 8, only our method shows a similar distribution to the real images.

Fig. 9 shows examples of brain slices generated using all comparison methods. Our method generates high-quality brain slices on three planes. Because of the random weight initialization, the comparison methods generate random images at the start of the training process without any meaningful patterns. Due to its training on sufficient 2D axial slices, our model exhibits good anatomy right at the beginning. Combining our Split&Shuffle design with this anatomy prior allows our model to generate better results with fewer parameters.

5. Conclusion

The purpose of our study was to address the important problem of generating reliable synthetic 3D medical images. The lack of annotated 3D data and inefficient parameter settings hinder the effective training of 3D medical generative models. A novel GAN model (i.e., 3D Split&Shuffle-GAN) is proposed to remedy these problems from two perspectives: training strategy and network architecture. For the training strategy, we used the weight inflation technique to pre-train a 2D GAN model and inflate the 2D convolution weights as a favorable method for initializing a 3D GAN model. For network architecture, we devised parameter-efficient Channel Split&Shuffle modules for the discriminator and generator of the GAN. We conducted comprehensive experiments to determine the best weight inflation variant and network architecture design. The effectiveness of our method is verified on both the heart and brain datasets. Further exploration of network weight initialization strategies beyond inflation and the design of new architectures will be completed in the future.

Funding

This work was supported by MRFF Frontier Health and Medical Research - RFRHPI000147. The computation in this work was partly supported by DUG facilities via Harry Perkins.

Declaration of Competing Interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Girish Dwivedi reports a relationship with Artrya Pty Ltd. that includes: consulting or advisory and equity or stocks.

References

- O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al., Imagenet large scale visual recognition challenge, Int. J. Comput. Vis. 115 (3) (2015) 211–252.
- [2] A. Krizhevsky, I. Sutskever, G.E. Hinton, Imagenet classification with deep convolutional neural networks, Adv. Neural Inf. Process. Syst. 25 (2012).
- [3] M. Frid-Adar, E. Klang, M. Amitai, J. Goldberger, H. Greenspan, Synthetic data augmentation using GAN for improved liver lesion classification, in: 2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018), IEEE, 2018, pp. 289–293.
- [4] Y. Zhou, Z. Yang, H. Zhang, E.I.-C. Chang, Y. Fan, Y. Xu, 3D Segmentation guided style-based generative adversarial networks for PET synthesis, IEEE Trans. Med. Imaging 41 (8) (2022) 2092–2104.
- [5] Y. Song, S. Ren, Y. Lu, X. Fu, K.K. Wong, Deep learning-based automatic segmentation of images in cardiac radiography: a promising challenge, Comput. Methods Programs Biomed. (2022) 106821.
- [6] X. Bian, X. Luo, C. Wang, W. Liu, X. Lin, Dda-net: unsupervised cross-modality medical image segmentation via dual domain adaptation, Comput. Methods Programs Biomed. 213 (2022) 106531.
- [7] E.H. Nguyen, H. Yang, R. Deng, Y. Lu, Z. Zhu, J.T. Roland, L. Lu, B.A. Landman, A.B. Fogo, Y. Huo, Circle representation for medical object detection, IEEE Trans. Med. Imaging 41 (3) (2022) 746–754.
- [8] Z. Wu, W. Liao, C. Yan, M. Zhao, G. Liu, N. Ma, X. Li, Deep learning based MRI reconstruction with transformer, Comput. Methods Programs Biomed. 233 (2023) 107452.
- [9] T. Makela, P. Clarysse, O. Sipila, N. Pauna, Q.C. Pham, T. Katila, I.E. Magnin, A review of cardiac image registration methods, IEEE Trans. Med. Imaging 21 (9) (2002) 1011–1021.
- [10] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, Adv. Neural Inf. Process. Syst. 27 (2014).
- [11] Q. Yang, P. Yan, Y. Zhang, H. Yu, Y. Shi, X. Mou, M.K. Kalra, Y. Zhang, L. Sun, G. Wang, Low-dose CT image denoising using a generative adversarial network with wasserstein distance and perceptual loss, IEEE Trans. Med. Imaging 37 (6) (2018) 1348–1357.
- [12] T.M. Quan, T. Nguyen-Duc, W.-K. Jeong, Compressed sensing MRI reconstruction using a generative adversarial network with a cyclic loss, IEEE Trans. Med. Imag. 37 (6) (2018) 1488–1497.
- [13] V. Kearney, B.P. Ziemer, A. Perry, T. Wang, J.W. Chan, L. Ma, O. Morin, S.S. Yom, T.D. Solberg, Attention-aware discrimination for MR-to-CT image translation using cycle-consistent generative adversarial networks, Radiology: Artif. Intell. 2 (2) (2020) e190027.
- [14] A. Ferreira, J. Li, K.L. Pomykala, J. Kleesiek, V. Alves, J. Egger, Gan-based generation of realistic 3D data: a systematic review and taxonomy, arXiv preprint arXiv:2207.01390 (2022).
- [15] P. Greenland, M.J. Blaha, M.J. Budoff, R. Erbel, K.E. Watson, Coronary calcium score and cardiovascular risk, J. Am. Coll. Cardiol. 72 (4) (2018) 434–447.
- [16] R. Gharleghi, A. Sowmya, S. Beier, Transient wall shear stress estimation in coronary bifurcations using convolutional neural networks, Comput. Methods Programs Biomed. 225 (2022) 107013.
- [17] F. Isensee, P. Kickingereder, W. Wick, M. Bendszus, K.H. Maier-Hein, Brain tumor segmentation and radiomics survival prediction: Contribution to the brats 2017 challenge, in: International MICCAI Brainlesion Workshop, Springer, 2017, pp. 287–297.
- [18] J. Cai, Z. He, Z. Zheng, Q. Xu, C. Hu, M. Huo, Learning global dependencies based on hierarchical full connection for brain tumor segmentation, Comput. Methods Programs Biomed. 221 (2022) 106925.
- [19] M. Monfort, A. Andonian, B. Zhou, K. Ramakrishnan, S.A. Bargal, T. Yan, L. Brown, Q. Fan, D. Gutfreund, C. Vondrick, et al., Moments in time dataset: one million videos for event understanding, IEEE Trans. Pattern Anal. Mach. Intell. 42 (2) (2019) 502–508.
- [20] J. Carreira, E. Noland, C. Hillier, A. Zisserman, A short note on the kinetics-700 human action dataset, arXiv preprint arXiv:1907.06987 (2019).
- [21] Stanford, Coca coronary calcium and chest cts dataset, 2022, https:// stanfordaimi.azurewebsites.net/datasets/e8ca74dc-8dd4-4340-815a-60b41f6cb2aa.
- [22] G. Kwon, C. Han, D.-s. Kim, Generation of 3D brain MRI using auto-encoding generative adversarial networks, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2019, pp. 118–126.
- [23] M.W. Weiner, D.P. Veitch, P.S. Aisen, L.A. Beckett, N.J. Cairns, R.C. Green, D. Harvey, C.R. Jack Jr, W. Jagust, J.C. Morris, et al., The Alzheimer's disease neuroimaging initiative 3: continued innovation for clinical trial improvement, Alzheimer's & Dementia 13 (5) (2017) 561–571.
- [24] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, T. Aila, Analyzing and improving the image quality of stylegan, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2020, pp. 8110–8119.
- [25] J. Carreira, A. Zisserman, Quo vadis, action recognition? A new model and the kinetics dataset, in: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017, pp. 6299–6308.

- [26] A.G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, H. Adam, Mobilenets: efficient convolutional neural networks for mobile vision applications, arXiv preprint arXiv:1704.04861 (2017).
- [27] T. Karras, S. Laine, T. Aila, A style-based generator architecture for generative adversarial networks, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2019, pp. 4401–4410.
- [28] D. Ravi, S.B. Blumberg, S. Ingala, F. Barkhof, D.C. Alexander, N.P. Oxtoby, A.D.N. Initiative, et al., Degenerative adversarial neuroimage nets for brain scan simulations: application in ageing and dementia, Med. Image Anal. 75 (2022) 102257.
- [29] E. Jung, M. Luna, S.H. Park, Conditional GAN with an attention-based generator and a 3D discriminator for 3D medical image generation, in: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VI 24, Springer, 2021, pp. 318–328.
- [30] S. Hu, W. Yu, Z. Chen, S. Wang, Medical image reconstruction using generative adversarial network for Alzheimer disease assessment with class-imbalance problem, in: 2020 IEEE 6th international conference on computer and communications (ICCC), IEEE, 2020, pp. 1323–1327.
- [31] S. Hu, J. Yuan, S. Wang, Cross-modality synthesis from MRI to PET using adversarial u-net with different normalization, in: 2019 international conference on medical imaging physics and engineering (ICMIPE), IEEE, 2019, pp. 1–5.
- [32] S. Hu, Y. Shen, S. Wang, B. Lei, Brain MR to PET synthesis via bidirectional generative adversarial network, in: Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part II 23, Springer, 2020, pp. 698–707.
- [33] S. You, B. Lei, S. Wang, C.K. Chui, A.C. Cheung, Y. Liu, M. Gan, G. Wu, Y. Shen, Fine perceptive GANs for brain MR image super-resolution in wavelet domain, IEEE Trans. Neural Netw. Learn. Syst. (2022).
- [34] B. Hu, C. Zhan, B. Tang, B. Wang, B. Lei, S.-Q. Wang, 3-D Brain reconstruction by hierarchical shape-perception network from a single incomplete image, IEEE Trans. Neural Netw. Learn. Syst. (2023).
- [35] S. Hong, R. Marinescu, A.V. Dalca, A.K. Bonkhoff, M. Bretzner, N.S. Rost, P. Golland, 3D-StyleGAN: A style-based generative adversarial network for generative modeling of three-dimensional medical images, in: DGM4MICCAI workshop in International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2021, pp. 24–34.
- [36] A. Arnab, M. Dehghani, G. Heigold, C. Sun, M. Lučić, C. Schmid, Vivit: a video vision transformer, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 6836–6846.
- [37] A. Novikov, D. Podoprikhin, A. Osokin, D.P. Vetrov, Tensorizing neural networks, Adv. Neural Inf. Process. Syst. 28 (2015).
- [38] V. Lebedev, Y. Ganin, M. Rakhuba, I. Oseledets, V. Lempitsky, Speeding-up convolutional neural networks using fine-tuned cp-decomposition, arXiv preprint arXiv:1412.6553 (2014).
- [39] J. Kossaifi, A. Toisoul, A. Bulat, Y. Panagakis, T.M. Hospedales, M. Pantic, Factorized higher-order CNNs with an application to spatio-temporal emotion esti-

mation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 6060–6069.

- [40] Y.-D. Kim, E. Park, S. Yoo, T. Choi, L. Yang, D. Shin, Compression of deep convolutional neural networks for fast and low power mobile applications, arXiv preprint arXiv:1511.06530 (2015).
- [41] K. Hara, H. Kataoka, Y. Satoh, Can spatiotemporal 3D CNNs retrace the history of 2D CNNs and Imagenet? in: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition, 2018, pp. 6546–6555.
- [42] F.N. landola, S. Han, M.W. Moskewicz, K. Ashraf, W.J. Dally, K. Keutzer, Squeezenet: alexnet-level accuracy with 50x fewer parameters and < 0.5 mb model size, arXiv preprint arXiv:1602.07360 (2016).
- [43] X. Zhang, X. Zhou, M. Lin, J. Sun, Shufflenet: an extremely efficient convolutional neural network for mobile devices, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 6848–6856.
 [44] O. Kopuklu, N. Kose, A. Gunduz, G. Rigoll, Resource efficient 3D convolutional
- [44] O. Kopuklu, N. Kose, A. Gunduz, G. Rigoll, Resource efficient 3D convolutional neural networks, in: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops, 2019. 0–0
- [45] T. Karras, M. Aittala, J. Hellsten, S. Laine, J. Lehtinen, T. Aila, Training generative adversarial networks with limited data, Adv. Neural Inf. Process. Syst. 33 (2020) 12104–12114.
- [46] K. Weiss, T.M. Khoshgoftaar, D. Wang, A survey of transfer learning, J. Big Data 3 (1) (2016) 1–40.
- [47] Y. Wang, A. Gonzalez-Garcia, D. Berga, L. Herranz, F.S. Khan, J.v.d. Weijer, Minegan: effective knowledge transfer from GANs to target domains with few images, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 9332–9341.
- [48] S. Chen, K. Ma, Y. Zheng, Med3d: transfer learning for 3d medical image analysis, arXiv preprint arXiv:1904.00625 (2019).
- [49] R. Solovyev, A.A. Kalinin, T. Gabruseva, 3D Convolutional neural networks for stalled brain capillary detection, Comput. Biol. Med. 141 (2022) 105089.
- [50] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 770–778.
- [51] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, S. Hochreiter, Gans trained by a two time-scale update rule converge to a local nash equilibrium, Adv. Neural Inf. Process. Syst. 30 (2017).
- [52] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, Z. Wojna, Rethinking the inception architecture for computer vision, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 2818–2826.
- [53] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, A.C. Courville, Improved training of wasserstein GANs, Adv. Neural Inf. Process. Syst. 30 (2017).
- [54] J. Wu, C. Zhang, T. Xue, B. Freeman, J. Tenenbaum, Learning a probabilistic latent space of object shapes via 3D generative-adversarial modeling, Adv. Neural Inf. Process. Syst. 29 (2016).
- [55] M. Rosca, B. Lakshminarayanan, D. Warde-Farley, S. Mohamed, Variational approaches for auto-encoding generative adversarial networks, arXiv preprint arXiv:1706.04987 (2017).