



# Use of Improved Generative Adversarial Network (GAN) Under Insufficient Data

Pallavi Adke<sup>1</sup>, Ajay Kumar Kushwaha<sup>2</sup>(✉), Supriya M. Khatavkar<sup>2</sup>,  
and Dipali Shende<sup>1</sup>

<sup>1</sup> Pimpri Chinchwad College of Engineering and Research, Pune, India

<sup>2</sup> Bharati Vidyapeeth (Deemed to be University) College of Engineering, Pune, India  
akkushwaha@bvucoep.edu.in

**Abstract.** The article covers enhancements to the generative adversarial network (GAN) model's architecture and training, enabling stable training in the absence of sufficient data. An improved generative adversarial network (GAN) architecture has been proposed. These improvements are then applied to the augmentation of a dataset on tyre joint defects, which is utilised for classification applications. The dataset used has a higher percentage of conformity images and is quite uneven. It is difficult to create precise defect classification models given this uneven and constrained dataset of defect identification. So, in the work that is being presented, research is done to expand the defect dataset and improve the balance between the various defect classifications. Indeed, the quality of generated images has considerably improved as a result of recent developments in generative adversarial networks (GANs). Deep learning models in the GAN class combine a generator network with a discriminator network. The current study reveals that the recommended augmented GAN model is useful in enhancing the performance classification model under a small dataset. The generated effects of progressed GAN are evaluated using the Fréchet Inception Distance (FID) score, which indicates extensive development over the styleGAN architecture. Additional dataset augmentation exams making use of generated photos monitor a 10% boom in category version precision in comparison to the preliminary dataset. To evaluate the effectiveness of GAN-generated picture augmentation, PCA plots can be used to visualize the distribution of real and augmented images in a lower-dimensional space.

**Keywords:** Generative Adversarial Network (GAN) · Data Augmentation · defect diagnosis · insufficient dataset

## 1 Introduction

Today, with little data, deep learning systems in the field of computer vision might suffer greatly. With an unbalanced dataset, the deep learning model's accuracy may further deteriorate. In the literature [1–5], it has been shown that, in the field of medical images, the availability of large datasets is very poor where it is needed to generate

images that are similar to the original images. The task of defect identification in an automated inspection process requires the model to locate defects in the input images and categorize them according to the type of defect. The procedure of gathering images for training such a model takes time because samples must be collected over time from the pertinent inspection line. Another drawback of this dataset is that it may contain a disproportionately high number of samples from the normal class. This is clear since every invention process is built to create samples that comply. Producing normal samples from the invention process to balance the dataset is extremely unrealistic and expensive.

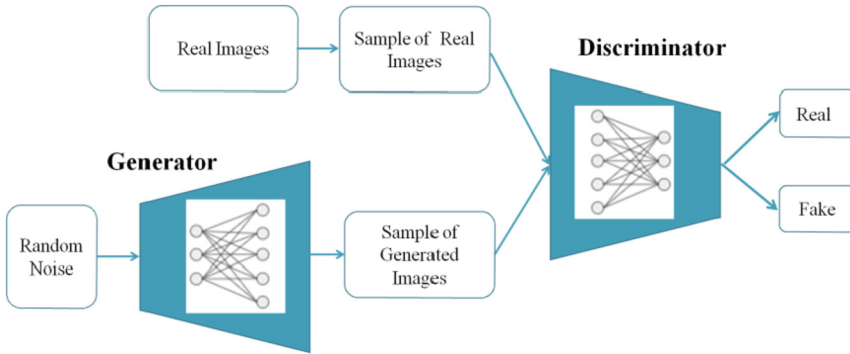
To improve the dataset that is now accessible, common picture augmentation techniques have been created. These methods change the original images using label-invariant and semantically preserved transformations. Such methods include focusing in and out, arbitrary flips, arbitrary shifts, rotations, and brightness changes, among others. Since most augmented images are only changes to existing images, their utility in capturing the entire probability distribution of an input dataset is limited. Furthermore, the use of these strategies depends on the problem. Generative Adversarial Networks (GANs) have been investigated as a potential solution to address data augmentation issues, especially in improving the precision of classification for recognizing defects. Conventional data augmentation techniques, such as image rotation, flipping, or scaling, have limitations in their ability to generate diverse and realistic variations of the original data. GANs offer a promising approach to overcome these shortcomings and generate synthetic data that can effectively enhance classification tasks. The primary training goal for GANs is to confine the distribution of actual data. The GAN characteristic is especially useful for jobs involving augmentation because GAN-generated images would completely cover the underlying distributions of real datasets. Additionally, it may result in less overfitting of the classification model.

To handle the low and unbalanced datasets at hand, the research effort reported in this paper explores current state-of-the-art developments in GAN. Changes to the GAN architecture, loss function, data augmentation, and regularization approaches are all included in these upgrades. The approach focuses on extracting tiny features from larger variation-generated images. For a small number of training images, this task is extremely difficult.

The structure of the paper is as follows: In Sect. 2, discuss the basic GAN model. The approaches utilized to enhance the underlying StyleGAN architecture and the specifics of studies using a suggested improved GAN to create augmented images are described in Sect. 3. In Sect. 4, GAN training uses regularizing strategies to boost convergence and stability. Experimental results are discussed in Sect. 5 and the conclusion is followed in Sect. 6.

## 2 Related Work

Generative Adversarial Networks (GANs) can generate sample images using the same distribution as the input real dataset (P data) [6]. GAN consists of a deep neural network model that is generally employed to produce artificial images in accordance with the distribution of training data. The GAN's basic architecture is depicted in Fig. 1.



**Fig. 1.** Basic GAN model

It includes the Generator and Discriminator models. The Generator model's primary goal is to produce samples that closely resemble the distribution of real data. The Discriminator, on the other hand, seeks to determine if the samples that are given to it are authentic or fraudulent.

In GAN, the generator receives a noise vector as an input to produce fresh samples. The generator network takes a random noise vector, typically drawn from a multivariate Gaussian distribution, and maps it to the space of the target data distribution. The typical random distribution used to generate this noise is: The Generator picks up how to correspond features in the output images to common noise. Convolutional neural networks are modelled as both Generators and Discriminators for image generation problems. The discriminator model performs down sampling operations that reduce the spatial dimensions of the input images. It then outputs a probability indicating the likelihood that the input is a real image from the original dataset. On the other hand, the generator model performs up sampling or up-convolution operations, also known as transpose convolution or deconvolution. It takes the noise vector as input and gradually increases the spatial dimensions of the data to generate synthetic images. The generator aims to create realistic enough images to fool the discriminator into classifying them as real. Generator and discriminator compete against one another in an adversarial setting during GAN Training. Over time, both the Generator and the Discriminator get more adept at producing samples that are difficult for the discriminator to discern between the real and defective images.

Numerous researchers have tried to use GAN for generating data since it was introduced in 2014. Aggarwal et al. [7] review GAN applications for pandemic and medical applications. The use of GAN for fake picture synthesis is shown to increase datasets, protect patient privacy, and lower additional costs associated with medical imaging procedures. GAN has been utilized by Gao et al. [8] to improve machine defect detection datasets. With GAN-generated datasets, they have shown gains in classifier accuracy. Ackey et al. employ GAN for anomaly detection [9]. Their methodology has produced a receiver operating characteristic curve area under the curve for recognizing abnormal/nonconforming samples that is 92 percent. To enhance labelled datasets for augmented reality applications, Ma et al. [10] investigated the 3D generating capabilities of

GAN. Researchers have investigated numerous intriguing applications of GAN in the fields of picture preprocessing, painting, super resolutions, image background domain change etc.

Researchers have conducted studies to better understand and improve the behavior of GAN training, with a focus on producing high-resolution images of higher quality. Researchers have explored techniques to generate high-resolution images using GANs. For example, in the case of the FFHQ dataset [11] and the LSUN automobile dataset, studies have achieved impressive results with FID (Fréchet Inception Distance) scores as low as 2.84 and 2.32, respectively. These low FID ratings show a high degree of visual quality and diversity similarity between the generated images and the genuine photos. Label conditioning was added to the StyleGAN architecture by Oeldorf et al. [12] during the picture generation process. Label conditioning gives users the ability to modify the generated images in accordance with extra data, such as class labels or categorical features. With this update, the created images are more diverse and controllable, allowing for finer-grained editing and control over particular image features. There have been various studies on regularizing methods in the area of GAN training stability. For trained GANs, Zhang et al. [13] presented a consistency regularization technique that focuses on improving the reliability of discriminator predictions for related images through semantic preservation augmentations. By encouraging the discriminator to provide consistent and reliable predictions for augmented versions of the same image, the GAN training becomes more stable. This regularization encourages the discriminators to focus on capturing image structural details, and it also helps to improve the flow of gradients from the discriminator to the generator, leading to better convergence. Mescheder [14] suggested a gradient penalty technique to address the issue of Lipschitz continuity in GANs. Lipschitz continuity ensures that small changes in the input space result in small changes in the output space, which aids in stable and reliable training. To enforce Lipschitz continuity, a gradient-based penalty is added to the discriminator's loss function. This penalty constrains the magnitude of gradients computed by the discriminator, ensuring smoother prediction landscapes. By imposing this constraint, the discriminator can take fewer gradient steps, leading to improved convergence and training stability. As a result, the mapping of latent vectors to picture features is smoother and more untangled. Numerous studies employ augmentation [15, 16] and regularization to address the problems with low training data. The methods section goes over these in more depth.

### 3 Methodology

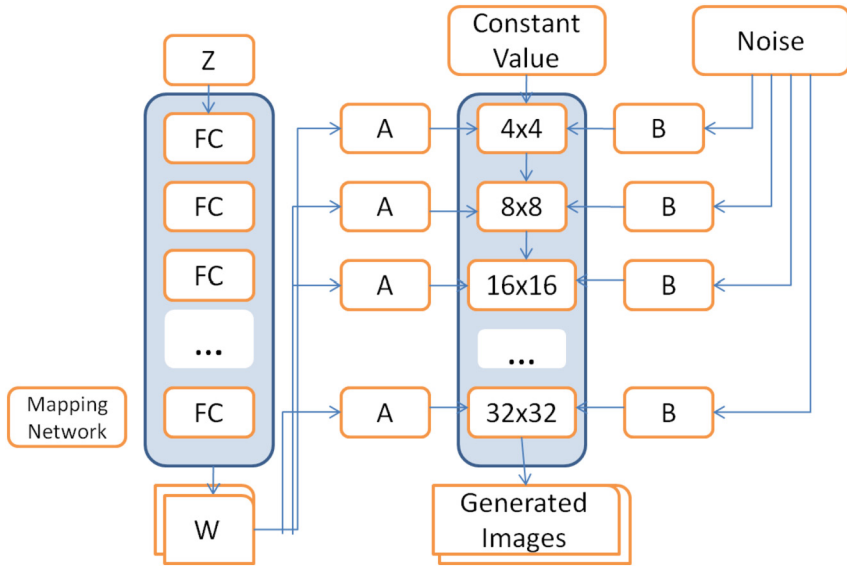
The primary goal of the work that is being presented is to create high-quality images of nonconformities that will aid in the subsequent process of image categorization. StyleGAN, developed by Karras et al. [17], is the foundation of the GAN architecture used for the current challenge. The current study incorporates the following GAN model and training enhancements:

#### 3.1 StyleGAN

Progressive GAN architecture is expanded upon by StyleGAN. The Generator's gradual expansion aids in the production of high-resolution images with better quality. To capture

small details in high-resolution images, one approach is to separate the training of low-level features from high-level features. This separation allows for more effective learning and representation of both low and high-level details in the generated images. The progressive network is added to by StyleGAN together with the mapping network. The input latent noise is transformed into intermediate vectors using the mapping network. As a result, entangled characteristics in generated images are reduced.

These transitional vectors are input into the Generator network at different stages to provide the user with additional control over the generated pictures. Adaptive Instance Normalization (AdaIN) layers are used for the injection in order to fit the generator feature map style to the input vector. The addition of random noise at each stage creates stochastic variation in the final images.



**Fig. 2.** StyleGAN model with a mapping network and a progressive Generator. The operations on layers “B” are noise scaling while layers “A” are affine transformation.

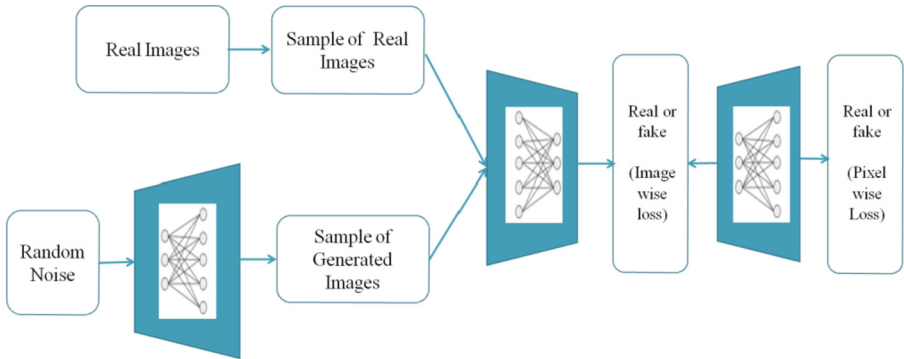
The Generator network is mirrored in the discriminator with ever-smaller images. By infusing distinct noise vectors into the generator at different phases, style mixing regularization is accomplished. Figure 2 depicts StyleGAN in broad strokes.

### 3.2 Discriminator for U-Net

The StyleGAN architecture’s discriminator categorizes the overall image as real or fake. To create locally coherent structures in images, the loss gradients that are so formed are only of limited use. A U-Net-based discriminator has been proposed by Schoenfeld et al. [18]. Figure 3 shows a block diagram of a U-Net-based GAN.

The U-net GAN may be trained using feedback at both the global and pixel levels. A decoder model of the discriminator gives information at the per-pixel level, whereas

an encoder model provides information at the global level of input images. As seen in our work, per-pixel content is important for creating images with conceptual relatedness based on true organization and capturing superior, nuanced information in images. Both complex and minute details of an image are transferred via skip links between the encoder and decoder models.



**Fig. 3.** Model U-net GAN

The U-net structure is added to the StyleGAN architecture created for the training. According to U-net GAN, StyleGAN's discriminator and loss algorithms were adjusted. The architecture's generator doesn't alter.

### 3.3 Data Augmentation in GAN Training

With insufficient training data, the image quality of GAN-generated images might substantially degrade. The discriminator may quickly overfit and stop delivering usable gradients to train the generator if it memorizes the key features from the training dataset. This results in bad-quality images and approaches collapse. Numerous studies have been conducted to use augmentation for training GAN, according to the literature [12]. The generator may create images that are both false and comparable to real photographs when conventional data augmentation is only used on real images. Unwanted distributions result in the samples that are created as a result. Instead, both real and fake images can be enhanced. This would provide a discriminator that is better at categorizing solely augmented images. Consequently, unconnected gradient flows following manipulations may prevent it from correctly identifying non-augmented generated images.

The application of differential augmentation offers a remedy for this [12]. Since all modifications made to both true and false images are differentiable, as the name suggests, gradients can proceed continuously from the discriminator to the generator. As a result, the generator continues to get precise training and the discriminator learns to recognize undamaged images from the appropriate target distribution. It is possible to differentiate between augmentations by utilizing the standard primary operations offered by deep learning frameworks.

The work by Karras et al. on various transformations that don't cause leakage in generated images can be found in [17]. They discovered that invertible transformations, such as pixel blitting, geometric, and colour transformations, have a better effect on produced images in terms of measuring metrics. These modifications are made with a nonzero probability (preferably lower than 0.8) during training to include non-augmented images as well.

### 3.4 Loss Functions

The current study's choice of loss function is mostly influenced by the presence of mode collapse in generated images. When the discriminator is overfitted to a small number of features in real picture distributions, mode collapse occurs. As a result, the generator often creates images that are solely appropriate for deceiving the discriminator based on those features. As a result, the generator is no longer able to make unique images. Mode collapse is more likely when there is less available data. Wasserstein loss with a gradient penalty is mostly used to address this issue (WGAN-GP) [19]. It teaches the discriminator to decrease the Wasserstein distance between the real distribution of the real samples ( $P_r$ ) and the generated distribution of the produced samples ( $P_g$ ). To ensure Lipschitz continuity close to the real data manifold, a consistency term is additionally attached to the WGAN-GP loss term. As described below, Wasserstein loss is applied in a non-saturating form.

Loss of the critic (discriminator):

$$L_{i \sim P_r}[D(i)] - L_{j \sim P_g}[D(G(j))] \tag{1}$$

where,  $P_r$  is distribution of real samples,

$P_g$  is distribution of produced samples,

$L_i$  is expected value of real inputs,

$D(i)$  is output of discriminator for real input ( $i$ ),

$G(j)$  is output of generator for random noise ( $j$ ),

$D(G(j))$  is discriminator output for input  $G(j)$ .

Generator loss:

$$L_{j \sim P_g}[D(G(j))] \tag{2}$$

where,  $L_j$  is expected value of generated data,

$D(G(j))$  is discriminator output for input  $G(j)$ .

The discriminator in WGAN-GP is known as the ‘‘critic’’ since it does not distinguish between real and fake images. A critic rates whether an image is authentic or fake. Here, the critic must adhere to 1-Lipschitz continuity to ensure that a loss assessed based on the output of the critic adheres to the Wasserstein distance metric. Lipschitz continuity is imposed by applying the gradient penalty, which is defined by the equation below, to reduce the norm of the gradients of the critic's output relative to an input to zero.

Gradient Penalty term:

$$GPT = L_{i \sim P_r, P_g} \left[ (\|\nabla_x D(i)\|_2 - 1)^2 \right] \tag{3}$$

where,  $L_j$  is expected value of generated data,  
 $\nabla_x D(i)$  is gradient of discriminator output of real data.  
 Term for consistency:

$$TFC = L_{i \sim Pr} \left[ (\|\nabla_x D(x)\|_2 - 1)^2 \right] \quad (4)$$

The formula for total critic loss is as follows:

$$L_{i \sim Pr}[D(i)] - L_{j \sim Pg}[D(G(j))] + \lambda * GPT + \lambda_1 * TFC \quad (5)$$

Scaling factors for the gradient penalty term and the consistency term, respectively, are in this case and 1. In the computation of critic loss, writers advise scaling the GPT term by a factor of 10 and the TFC term by 2.

## 4 Regularization

GAN training uses regularizing strategies to boost convergence and stability. These techniques can be classified according to how they implement network weights, gradients, and layer outputs. On the discriminator, most regularizing approaches are used. On generator weights, very few approaches, such as perceptual path length regularization, are used. The current research focuses on regularizing the discriminator to improve mode collapse and training stability. The discriminator is subjected to consistency regularization to impose equivariant behaviour for differential augmentation. CutMix-augmented pictures are used to put it into effect. Crops of actual and synthetic photos are mixed to supply those graphics. Consistency loss guarantees that a discriminator prediction for a CutMix photo and a group of predictions from its impartial vegetation fluctuate little or nothing from one another. Gradient penalty terms, mentioned within the preceding phase and as utilized in loss evaluations, even have a regularizing impact with the aid of maintaining gradient beneath cohesion and using Lipschitz continuity. The generator weights' exponential weight averaging tune is saved at some point in training. These averaging weights are implemented while developing photos for augmentation. Since common weights are much less sensitive to noise and outlier iterations at some stage in training, they offer photos of better quality.

The enhancements outlined above are used in the current work on picture augmentation utilizing GAN generation to create higher-quality images. To capture pixel-level information, a discriminator from a styleGAN model is converted to a U-NET architecture. To solve the lack of training dataset availability, differential augmentation is used. In order to lessen the problem of mode collapse and produce images with more variations, a better WGAN-GP loss term is utilized. By including a consistent loss term and a gradient penalty term in loss evaluations, a regularization effect is produced. Finally, images for augmentation are produced using the generator with exponentially moving average weights. The remainder of the essay will refer to this adapted GAN architecture as improved GAN.

## 5 Discussion of Experiments and Results

The tyre texture image recognition dataset [20] is used for training and testing. The dataset consists of 1028 images of cracked (oxidized) and normal images tyres. The dataset is divided into training and testing images. Using a dataset on tyre joint compliance, the proposed improved GAN's applicability is assessed. Images are produced by combining cracked (oxidized) and normal images from the dataset in a number of tests. Three methods are used to supplement the data. Table 1 provides a summary of all methods used for image production. Each defect's GAN model is trained using the single noise vector approach. The style merging method is used to train both distinct GAN models for each defect and normal image as well as a single GAN model for all defective photos. Whereas the latent interpolation technique is used for training all GAN models.

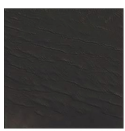
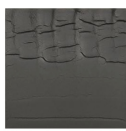
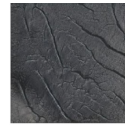

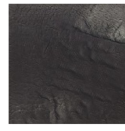

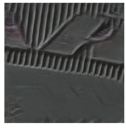

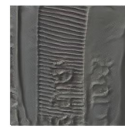





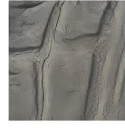

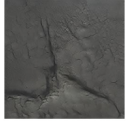
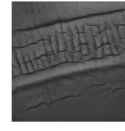
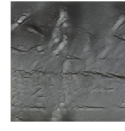
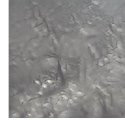
For each defect category, a separate GAN model is trained in the first method. Then, independent augmented images of each defect are produced using these trained models. In the second approach, a GAN model is trained using photos from all categories. Style merging on the trained generator creates enhanced visuals. Two separate nonconforming images' latent vectors are input into the styleGAN generator at various resolutions. This method of injecting style creates visuals that shift from one defect to another. As a result, we can create a dataset that allows us to change an image's defect category. The third technique develops a new GAN model using a collection of conforming and nonconforming photos from a single category. This demonstration can be utilized to present the imperfections with which it was prepared into a standard picture by utilizing fashion combining. The momentary and third strategies of information increase, which also include inactive vector addition to move pictures from one category to another, were created using Python 3.6 and the TensorFlow 2.1.0 system. Table 1 shows the pictures made using the proposed GAN.

The Microsoft Azure Machine Learning Services are used to train all models. One NVIDIA Tesla K80 GPU is used for computation. The final image resolution is 256 x 256 pixels. The quality of the generated images is evaluated using the Fréchet inception distance (FID) [21, 22]. A classification model that has been trained to categorize images into each defect or conformity (OK) category is used to assess the efficacy of augmentation. Convolutional neural networks are the foundation of the categorization model.

The basic StyleGAN model architecture is contrasted with the improved GAN model that is proposed. Utilizing FID, their performance is assessed. It has been noticed that improved image quality and variety are correlated with lower FID scores. The identical tyre joint defect datasets are used to train both designs and the results are compared. Their difference is shown in Table 2.

The FID score for improved GAN has significantly increased when compared to the styleGAN model, according to these data. The outcomes further demonstrate the value of improved GAN in enhancing generation quality with a constrained quantity of training images. Changes made to Improved GAN's architecture and training have improved the outcomes. In order to deal with limited dataset regimes, differential augmentation and consistency regularization have been implemented. Additionally, it maintains training for improved convergence. The UNET discriminator offers pixel-by-pixel feedback that

**Table 1.** GAN generated images for different defects in tyres.

Clusters	GAN Image_1	GAN Image_2	GAN Image_3	GAN Image_4	GAN Image_5
Cluster_1					
Cluster_2					
Cluster_3					
Cluster_4					

**Table 2.** Comparison of StyleGAN's performance with that of the planned Improved GAN FID Scores GAN Architecture

GAN Architecture	FID Scores		
	Defects 1	Defects 2	Defects 3
StyleGAN	162.9	159	158.1
Proposed Improved GAN	95.6	91.9	94.8

enhances the generated image quality and lowers the FID score. By reducing outlier noisy iterations and averaging the generator weights' exponential weights, the training oscillations and FID score are further decreased.

To investigate the impact of augmentation, the classifier model is initially trained on just real photos without incorporating any GAN-generated images. For classifier model training, all tests use the same conventional augmentations, such as horizontal flip, crop, and translate. To test the classifier model, only genuine photographs that were arbitrarily chosen from the original dataset are used. 10% of the real photos are used for testing, with the remaining 90% being used for training and validation. The accuracy of the trained classifier model is compared with other augmentation studies. Precision on a test dataset

is assessed and appears to be the cruellest of all test tests across all classes. Table 2 gives an outline of the results of all tests utilizing both created and genuine photographs. The results displayed here are the result of numerous classification models prepared on the same dataset to decrease instability. Accuracy on a test dataset is surveyed, and it shows up to be the cruellest of all test scores across all classes. Table 2 gives a diagram of the results of all tests utilizing both conveyed and verifiable photos. The results shown here are the result of various classification models arranged on the same dataset to decrease insecurity.

The dataset for this study was created using data from two stages of a manufacturing line. In the initial phase, 1183 samples were collected altogether. An additional 1108 samples were collected in the second stage, bringing the total to 2291. GAN models are trained using real data from the first stage, and reinforcement comes from generated images. Afterward, all genuine photographs from both stages are utilized for GAN preparation. The proficiency of the increase is surveyed freely for each set of created pictures from the two stages.

## 6 Conclusion

Incorporating recent innovations into GAN models for improved generated image quality is covered in the paper. Compared to styleGAN, the proposed improved GAN architecture yields significantly lower FID scores, which suggests improved image quality and generation variation. The paper discusses some architectural and training enhancements that can help GAN training converge more smoothly. Therefore, the improved GAN that has been proposed may produce a variety of images with fine features. Improved GAN is very helpful when supplementing small, unbalanced datasets. The accuracy of the picture classification tasks that follow has significantly improved thanks to an upgraded, balanced dataset. Experimental evidence from the principal components of the enhanced dataset supports the claim that generated images from the improved GAN can aid in improving the differentiation between various classification classes.

Due to limitations in computing power and processing speed, the experiments presented in this paper were only able to produce images that were 256 x 256 Pixels in size. In the case of smaller datasets, the effectiveness of picture augmentation by GAN produced images is high. Its applicability to huge datasets requires additional research. The present work's future scope includes using style-merged images for augmentation and adding GAN model enhancements with styleGAN2 architecture. For classes with lower classification recall, class wise augmentation can be explored.

## References

1. Kulkarni, P., Madathil, D.: A review of echocardiographic image segmentation techniques for left ventricular study. *ARPN J. Eng. Appl. Sci.* **13**(10), 3536–3541 (2018)
2. Kulkarni, P., Madathil, D.: A review on echocardiographic image speckle reduction filters. *Biomed. Res.* **29**, 12 (2018)
3. Kulkarni, P., Madathil, D.: Adaptive thresholding method for speckle reduction of echocardiographic images. *IETE J. Res.* **68**(2), 1034–1042 (2022)

4. Kulkarni, P., Madathil, D.: Fully automatic segmentation of LV from echocardiography images and calculation of ejection fraction using deep learning. *IJBET* **40**(3), 241 (2022)
5. Kulkarni, P., Madathil, D.: Echocardiography image segmentation using semi-automatic numerical optimisation method based on wavelet decomposition thresholding. *Int. J. Imaging Syst. Tech.* **31**(4), 2295–2304 (2021)
6. Goodfellow, I.J., et al.: Generative adversarial networks (2014)
7. Aggarwal, A., Mittal, M., Battineni, G.: Generative adversarial network: an overview of theory and applications. *Int. J. Inform. Manage. Data Insights* **1**, 100004 (2021)
8. Gao, X., Deng, F., Yue, X.: Data augmentation in fault diagnosis based on the Wasserstein generative adversarial network with gradient penalty. *Neurocomputing* **396**, 487–494 (2020)
9. Akcay, S., Atapour-Abarghouei, A., Breckon, T.P.: GANomaly: semi-supervised anomaly detection via adversarial training. In: Jawahar, C.V., Li, H., Mori, G., Schindler, K. (eds.) *ACCV 2018*. LNCS, vol. 11363, pp. 622–637. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-20893-6\\_39](https://doi.org/10.1007/978-3-030-20893-6_39)
10. Ma, Q., et al.: Learning to dress 3D people in generative clothing. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Seattle, WA, USA, pp. 6468–6477. IEEE (2020)
11. Kramberger, T., Potočník, B.: LSUN-Stanford car dataset: enhancing large-scale car image datasets using deep learning for usage in GAN training. *Appl. Sci.* **10**(14), 4913 (2020)
12. Oeldorf, C., Spanakis, G.: LoGANv2: Conditional style-based logo generation with generative adversarial networks. In: *2019 18th IEEE International Conference on Machine Learning and Applications (ICMLA)*, Boca Raton, FL, USA, pp. 462–468. IEEE (2019)
13. Zhang, H., Zhang, Z., Odena, A., Lee, H.: Consistency regularization for generative adversarial networks. [arXiv:1910.12027](https://arxiv.org/abs/1910.12027) (2020)
14. Mescheder, L., Geiger, A., Nowozin, S.: Which training methods for GANs do actually converge. [arXiv:1801.04406](https://arxiv.org/abs/1801.04406) (2020)
15. Zhao, S., Liu, Z., Lin, J., Zhu, J., Han, S.: Differentiable augmentation for data-efficient GAN training. [arXiv:2006.10738](https://arxiv.org/abs/2006.10738) (2020)
16. Sinha, A., Ayush, K., Song, J., Uzket, B., Jin, H., Ermon, S.: Negative data augmentation. [arXiv:2102.05113](https://arxiv.org/abs/2102.05113) (2021)
17. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4396–4405 (2019)
18. Schonfeld, E., Schiele, B., Khoreva, A.: A U-net based discriminator for generative adversarial networks. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8204–8213 (2020)
19. Wei, X., Gong, B., Liu, Z., Lu, W., Wang, L.: Improving the improved training of Wasserstein GANs: a consistency term and its dual effect (2018)
20. Siegel, J.: Oxidized and non-oxidized tire sidewall and tread images, Harvard Dataverse (2021)
21. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: GANs trained by a two-timescale update rule converge to a local Nash Equilibrium. In: *NeurIPS Proceedings*, vol. 30 (2017)
22. Kushwaha, A.K., Khatavkar, S.M., Biradar, D.M., Chougule, P.A.: Depth estimation and navigation route planning for mobile robots based on stereo camera. *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering* **472**, 180–191 (2023). [https://doi.org/10.1007/978-3-031-28975-0\\_15](https://doi.org/10.1007/978-3-031-28975-0_15)