International Journal of Innovative Research in Computer Science and Technology (IJIRCST)
ISSN (Online): 2347-5552, Volume-14, Issue-1, January 2026
DOI: https://doi.org/10.55524/ijircst.2026.14.1.4
Article ID IRP-1705, Pages 25-33
www.ijircst.org

Generation of Human Faces Using Generative Adversarial Network

Supritha P O[1], Prajwal[2], Saikumar Laxman Pujari[3], Siddartha R[4], and Adarsh Shendage[5]

[1]Assistant Professor, Department of Computer Science & Engineering, Sri Dharmasthala Manjunatheshwara Institute of Technology, Ujire, Karnataka, India

[2,3,4,5] BE Scholar, Department of Computer Science & Engineering, Sri Dharmasthala Manjunatheshwara Institute of Technology, Ujire, Karnataka, India

Correspondence should be addressed to Supritha P O; suprithapoov95@gmail.com

Received: 22 November 2025 Revised: 9 December 2025 Accepted: 23 December 2025

ABSTRACT- GANs have proven to be a powerful deep-learning framework for creating realistic synthetic images, finding wide use across various tasks in computer vision. In this work, we introduce a GAN-driven method for the generation of human face images using a deep convolutional structure. We trained our model on 9,090 resized facial images to 128 × 128 pixels from the without-mask portion of the Face Mask Lite Dataset. In this adversarial setup, the Generator synthesizes facial images from random noise vectors, while the Discriminator distinguishes between real and generated samples.

The main goal here is to build a lightweight, computation-friendly GAN that still yields visually convincing face images without resorting to heavy architectures. These experiments show that the model captures the essential facial features-symmetry, texture, and overall appearance-and generates a diverse set of synthetic faces. The evaluation combines the qualitative visual inspection of generated samples with quantitative analysis of Generator and Discriminator loss trends. The results point toward stable training, realistic face generation, and a preference for architectural simplicity.

KEYWORDS- Generative Adversarial Network, Fake Human Faces, Deep Learning, Generator, Discriminator.

I. INTRODUCTION

GANs are a class of deep learning models designed to generate synthetic data, particularly images. The core idea of GANs is based on adversarial learning, where two neural networks- the Generator and the Discriminator-are trained simultaneously. The Generator attempts to create realistic samples from random noise, while the Discriminator learns to distinguish real samples from generated ones. In this competitive loop, both improve over time through a process that generates ever more convincing outputs.

This has led to significant research interest in the field of human face generation, with broad applications including virtual avatars, animation, games, digital content creation, and the generation of privacy-preserving data, among many others. Unlike traditional image generation approaches that rely on handcrafted features, GANs automatically learns facial representations directly from training data. This enables the model to capture intricate facial attributes such as symmetry, skin texture, lighting variations, and structural consistency, enabling the creation of realistic faces that are nonexistent in reality.

Although advanced GANs like StyleGAN and its variants can provide highly photorealistic faces, they often require heavy computational resources and extensive training times. The cost of this might be prohibiting for academic purposes or generally for those that don't have a good hardware basis. This motivates the search for simpler, more efficient GAN architectures which do not sacrifice convincing visuals. In this work, we present a deep convolutional GAN, targeting simplicity, stability during the training process, and computational efficiency when generating human face images at resolution 128 × 128. Training and evaluation rely on the Face Mask Lite Dataset (no-mask subset).

Beyond performance, another important aspect is the ethical dimensions of realistic face generation. As the realism of AI-synthesized facial imagery improves, thoughtful design and vigilant evaluation are required to prevent potential abuse. The current research seeks to develop and evaluate a GAN-based face-generation system while providing insight into its efficacy, limitations, and responsible use.

II. RELATED WORK

Several studies have explored the application of Generative Adversarial Networks for human face generation and related image synthesis tasks. This section reviews key contributions that have influenced the development of GAN-based face generation models. Early research focused on learning overall facial structures from large-scale datasets using basic GAN architectures. Later works introduced deep convolutional models to improve training stability and capture finer facial details. More recent studies have examined high-resolution face synthesis and architectural improvements to address challenges such as mode collapse and training instability, while also emphasizing ethical considerations in synthetic face generation.

Sakshi Singh et al. [1] studied presents a GAN model trained on the CelebA dataset to generate realistic human faces. The system uses convolutional layers with batch normalization and LeakyReLU activation to learn detailed facial structures. It achieves stable training and produces clear, high-quality synthetic images that resemble real faces. The work demonstrates the effectiveness of DCGAN-style architectures for capturing complex facial patterns.

Md. Mahiuddin and Azad Chowdary [2] introduces a GAN system designed to generate lifelike faces to assist police investigations. The model can create realistic suspect-like faces from descriptive or incomplete inputs. The authors highlight concerns related to privacy, identity misuse, and ethical handling of synthetic facial data. Their findings show that GAN-generated faces can support forensic workflows when used responsibly

Vamsi Sai Krishna Katta et al. [3] proposed a hybrid model where DCGAN first generated low-resolution faces. The images are then enhanced using ESRGAN to improve sharpness and perceptual quality. The approach results in clearer textures, improved structure, and more realistic face outputs. Their method outperforms standard GAN models in detail preservation and visual accuracy.

Xin Wang et al. [4] reviewed more than 100 studies focused on detecting GAN-generated synthetic faces. Their analysis highlights the increasing difficulty of distinguishing real and fake faces due to model advancements. They also discuss challenges posed by adversarial attacks and human limitations in detecting fakes. The study emphasizes the urgent need for stronger and more reliable deep fakes detection mechanisms.

Ravinder Reddy and Raman [5] examines the use of Conditional GANS (CGANs) for editing facial features such as age, hairstyle, and expression. CGANs allow multiple variations of the same identity to be generated with controlled attributes. The authors discuss challenges in maintaining realism and diversity when modifying multiple attributes. Their findings show that CGANs provide flexible and personalized face-generation capabilities.

Goodfellow et al. [6] introduces the GAN framework, where a generator and discriminator compete in an adversarial process. The model learns data distributions without supervision and can create realistic synthetic images. It demonstrates the potential of adversarial learning for producing high-quality visual outputs. This paper forms the base architecture for nearly all later face-generation research.

Radford, Metz & Chintala [7] proposed a stable GAN architecture using convolutional layers and batch normalization. The model generates realistic face images and learns meaningful visual features from datasets like CelebA. Its design guidelines improved GAN training stability and output quality. DCGAN became a widely used baseline for face-generation applications.

Arjovsky et al. [8] model reduces mode collapse and provides a smoother learning signal for the generator. It leads to more consistent and reliable synthetic face generation. This work significantly influenced improvements in GAN training methods.

Gulrajani et al. [9] introduces a gradient penalty to enforce Lipschitz constraints during training. The method stabilizes GAN optimization and improves image realism. Its simple regularization greatly reduces common GAN training failures. WGAN-GP became a standard loss approach in face-generation research.

Miyato et al. [10] presents spectral normalization to control the discriminator's weight scaling. It stabilizes GAN training and prevents gradient explosions. The technique improves the quality and consistency of generated face images. Spectral normalization is widely used in modern GAN architectures for reliability.

Karras et al. [11] model trains GANs by gradually increasing image resolution during training. Progressive growth improves stability and enables extremely high-resolution face synthesis. The method produces detailed and photorealistic facial images. It marked a major step forward in generating high-quality face outputs

Karras, Laine & Aila [12], introduces a Style-Based Generator Architecture (StyleGAN). StyleGAN style modulation to control different levels of facial features. It generates highly realistic human faces with strong attribute separation. The architecture supports intuitive edits such as age, pose, and expression changes. StyleGAN became a landmark model in photorealistic face generation.

Karras et al. [13] improved Image Quality and Stability.StyleGAN2 corrects artifacts found in the original StyleGAN and improves detail sharpness. The architecture produces more accurate textures and smoother facial structures. It enhances identity preservation during facial image manipulation. StyleGAN2 is widely used for professional-grade synthetic face creation.

Isola et al. [14] proposed Pix2Pix uses paired datasets to translate structured inputs into realistic images. In face generation, it converts sketches or edge maps into full facial outputs. The model learns detailed pixel-level correspondences between domains. Its conditional framework is widely used for controlled face synthesis tasks

Zhu et al. [15] proposed CycleGAN enables image translation without paired training data using cycle consistency. It supports facial transformations such as aging, makeup transfer, and emotion changes. The method preserves identity while modifying style-related features. It is useful when paired facial datasets are unavailable.

Karras et al. [16] proposed StyleGAN2-ADA for Limited Data Training.StyleGAN2-ADA introduces adaptive data augmentation to avoid discriminator overfitting. It enables training high-quality GANs even when datasets are small. The method maintains visual fidelity without requiring large image collections. It is especially useful for customized face-generation tasks with restricted data.

Choi et al. [17] proposed StarGAN: Multi-Domain Facial Attribute Translation. StarGAN allows multiple attribute transformations using a single unified GAN model. It enables changes like gender, hair color, or facial expression within one network. The approach simplifies multi-domain training by avoiding separate models. StarGAN is widely used for flexible and efficient face editing.

Karras et al. [18] introduces StyleGAN3: Alias-Free Generative Adversarial Networks.StyleGAN3 removes aliasing artifacts and improves geometric consistency in generated faces. The model produces smoother transitions in animations and poses changes. It offers more stable synthesis for video and motion-based applications. This version enhances the realism and coherence of face-generation outputs.

He et al. [19] introduces AttGAN: Facial Attribute Editing. AttGAN edits specific face attributes while preserving all non-target details. The model uses reconstruction and classification constraints for accurate results. It generates natural-looking faces even after multiple attribute changes. AttGAN is effective for identity-preserving face manipulation tasks.

Härkönen et al. [20] introduces GANSpace: Interpretable GAN Latent Directions. GANSpace identifies meaningful latent directions such as age, pose, or lighting from pretrained GANs. These directions allow users to edit faces without retraining the model. The method provides intuitive and interactive control over face attributes. It is widely used for GAN-based face editing and visualization tools.

III. PROPOSED METHODOLOGY

A GAN-based system is presented for generating realistic human facial images from random noise inputs. Our model finds a good balance between a simple network architecture, model training stability, and visual realism, and this makes our model feasible to implement using a machine with lower processing capacity. Our model comprises two interlinked components: the Generator and the Discriminator.

In terms of processing input data, face images from the Face Mask Lite Dataset are resized uniformly to 128 × 128 pixels. Here, OpenCV is used for image processing. Since OpenCV reads images in BGR format by default, the images are converted to RGB format for accurate color representation. Subsequently, all pixels are normalized in the range [−1, 1] in order to be in accordance with the tanh activation function in the final output layer of the Generator.

The Generator receives a 100-dimensional vector sampled from a standard normal distribution. Such a vector contains random noise used for diversification of generated faces. Diverging from a traditional way of processing small feature maps, our model projects this vector onto a full image representation and combines both convolutional and transposed convolutional layers in a hybrid path in order to capture both a face in general and small particular details.

The Discriminator is a binary classifier used to distinguish real images from generated samples. The Discriminator takes in 128x128x3 RGB images and undergoes a series of convolutional layers with an increasing depth of features and down-sampling. Through this hierarchical feature extraction technique, it is able to detect edges, textures, and structural information in an image. The output of the Discriminator will be a probability representing how real an image is. The learning algorithm is based on adversarial optimization, where the Generator network and Discriminator network are updated in a two-step alternating manner. Every iteration, Discriminator learns from a batch of real images along with a batch of fake images produced from random noise. The real images are labeled as genuine, and generated images are labeled as not genuine. Based on this information from Discriminator, Generator trains to model images resembled with real faces. The Binary Cross-Entropy loss function is used to optimize both the Generator and the Discriminator. An RMSProp optimizer with controlled gradient updates is applied. With each epoch, Generator learns to model better the real facial feature distribution of a human face and generates diversified faces with random noise.

IV. DATASET DESCRIPTION

To train the proposed GAN model, the without-mask subset of the Face Mask Lite Dataset is used. The subset contains facial images where there are no masks covering the faces, thus making it perfect for learning facial traits such as facial shape and facial appearance. A total of 9,090 images were chosen for model training.

The dataset is very diverse in terms of head pose, lighting, background, and facial expressions. Such diversity is important in learning a robust generative model, which will prevent overfitting and allow a Generator to understand a wide variety of facial characteristics. Therefore, this model is ideal for generating a variety of realistic synthetic faces.

A standard preprocessing step is performed before the training takes place. The images are read using OpenCV in BGR format but are immediately changed to RGB format for accurate color display. Every image is resized to a standard resolution of 128x128 pixels.

To effectively train a GAN, pixels are normalized to [-1,1] intervals. Additionally, image arrays are preprocessed into numerical form and presented in mini-batch sizes of 32 for learning. Such a learning framework ensures equity among samples and promotes efficient learning. Normalization is performed based on a tanh activation function in the final layer of a Generator in a Generator network during stable learning in a GAN mode.

Figure 1 shows sample images from the Face Mask Lite Dataset (without-mask subset), which were used to train the proposed GAN and provide variations in lighting, pose, and facial appearance for improved learning.

Figure 1: Face Mask Lite Dataset

V. MODEL ARCHITECTURE

The proposed Generative Adversarial Network comprises two deep learning models, referred to as the Generator and the Discriminator. The two models are trained simultaneously in an adversarial manner, where the Generator produces synthetic face images and the Discriminator distinguishes real images from generated ones. The rivalry between these two models leads to learning of the face feature distribution in the face dataset. In the below Figure 2 illustrates the overall GAN workflow used in work, highlighting how the Generator and training.

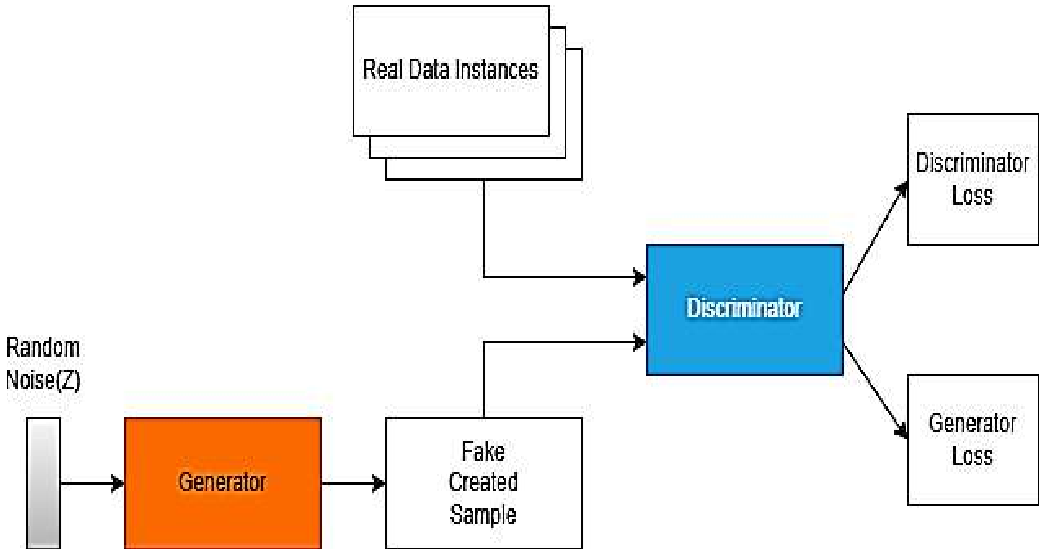

Figure 2: GAN architecture

A. *Generator*

The Generator acts as the "engine" responsible for generating the synthetic images of human faces with a resolution of 128x128x3 pixels. This Generator begins with a 100-dimensional vector obtained from a standard normal distribution, and this serves to introduce randomness into the results. The vector is first passed through a dense layer and subsequent to this, a spatial feature map with a "blueprint" of a face is obtained.

Further, this model combines convolutional layers and transposed convolutional layers to refine this image. Here, convolutional layers are responsible for learning features at various levels, including facial geometry, from an image, and transposed convolutional layers are used to increase the resolution in each image systematically by a fixed amount in order to recover facial detail.

To improve the stability of training, batch normalization layers are applied at multiple stages to reduce internal covariate shift. Leaky ReLUs are used in all the intermediate layers to allow non-linearity without disrupting the flow of gradients. The final output layer consists of a transposed convolution with a tanh activation function, which gives an output of an RGB image with pixels in the range −1 to 1.

Figure 3 presents the internal architecture of the Generator, showing how the latent vector is gradually transformed into a full-resolution synthetic facial image.

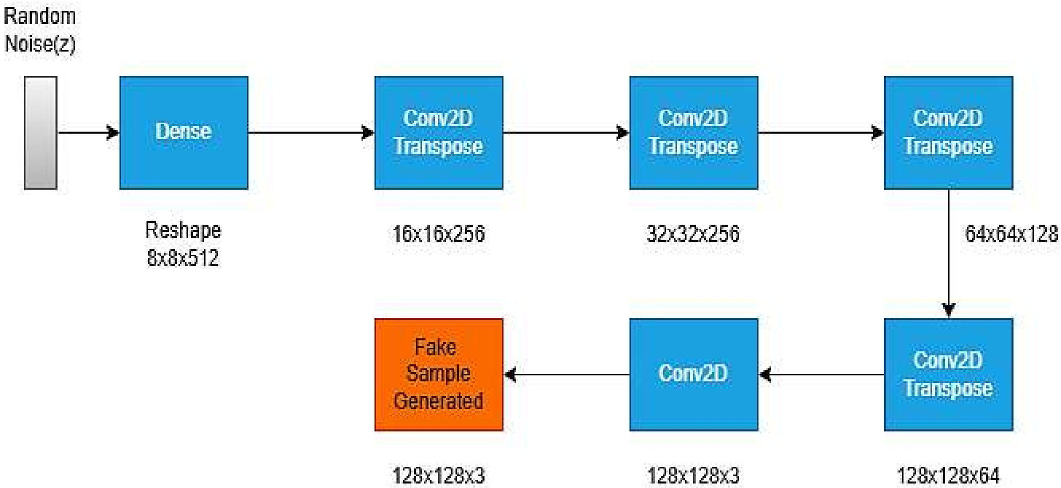

Figure 3: Generator architecture

B. *Discriminator*

The Discriminator is a binary classifier that determines whether a given image is real or fake, created by the Generator. The Discriminator receives a 128x128x3 RGB

image, which can be either a picture from a dataset or a product of the Generator.
The architecture comprises a series of convolutional layers with increasing numbers of features, which are achieved with strides to reduce image dimensions. With each reduction in image dimensions, the model learns increasingly abstract representations of facial characteristics such as edges, textures, and structures. Every convolutional layer incorporates a Leaky ReLU activation function to prevent dead regions in neurons. Additionally, batch normalization is performed in the middle layers for better stability during training.
Once the final convolutional layer is achieved, the feature maps are flattened and passed to a fully connected layer with sigmoid activation. As a result, a scalar output in the range [0, 1] is obtained, representing how likely the input image is to be real. The output of Discriminator acts as an adversarial signal for Generator.

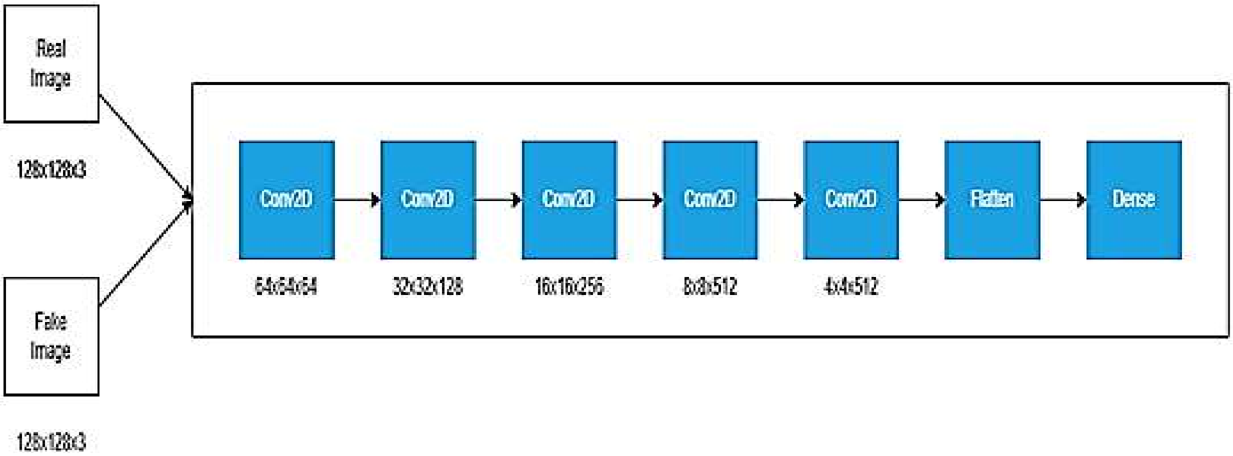

Figure 4: Discriminator architecture

In the above Figure 4 displays the Discriminator architecture, which progressively extracts hierarchical facial features to classify images as real or synthetic.

VI. TRAINING DETAILS

The training of the proposed GAN model follows an adversarial optimization strategy in which the Generator and Discriminator are trained alternately. The goal of this process is to enable the Generator to produce realistic human face images, while the Discriminator learns to accurately distinguish real images from generated ones.
Training was carried out in the Kaggle notebook environment using GPU acceleration to meet computational requirements efficiently. The dataset was divided into mini-batches of sizes 32, and each batch was used to perform one adversarial training step. Random noise vectors of dimension 100 were sampled from a standard normal distribution and provided as input to the Generator.

A. Adversarial Learning Objective

The overall objective of a Generative Adversarial Network can be expressed as a minimax game between the Generator Gand the Discriminator D, defined as:

$$\min_G \max_D V(D, G = \mathbb{E}_{x\sim p_{data}(x)}[\log D(x)] + \mathbb{E}_{z\sim p_z(z)}[\log(1 - D(G(z)))]$$

where x represents real images sampled from the training dataset, z denotes a random noise vector sampled from a normal distribution, G(z)is the generated image, and D(.)represents the probability output of the Discriminator.

B. Discriminator Loss Function

The Discriminator is trained to correctly classify real images as real and generated images as fake. The loss function used for training the Discriminator is Binary Cross-Entropy loss and is defined as:

$$L_D = -\mathbb{E}_x[\log D(x)] - \mathbb{E}_z[\log(1 - D(G(z)))]$$

This loss penalizes the Discriminator when it incorrectly classifies real images or fails to identify generated images as fake. During each training iteration, the Discriminator is updated using both real and generated image samples.

C. Generator Loss Function

The Generator is optimized to produce images that can successfully fool the Discriminator. The loss function used for training the Generator is given by:

$$L_G = -\mathbb{E}_z[\log D(G(z))]$$

This loss encourages the Generator to increase the likelihood that generated images are classified as real by the Discriminator. As training progresses, minimizing this loss helps improve the realism of generated facial images.

D. Optimization Strategy

Both the Generator and the Discriminator were optimized using the RMSprop optimizer with a learning rate of 0.0001. Gradient clipping was applied to stabilize training and prevent sudden parameter updates. Binary Cross-Entropy loss was employed consistently across both networks, aligning with the sigmoid activation used in the final layer of the Discriminator.
Training was conducted for multiple epochs, and the loss values of both networks were recorded at each epoch to monitor convergence and training stability. This adversarial learning strategy enabled the Generator to gradually learn meaningful facial representations and synthesize visually realistic human face images.

VII. RESULTS AND DISCUSSION

The effectiveness of the proposed GAN model is evaluated based on the learning behaviors of the Generator and the Discriminator models and by inspecting the generated facial images qualitatively since the end goal for generative models will always be to produce realistic results from the start—visually at least.

A. Training Loss Analysis

In below Figure 5 shows the training losses of the Generator and the Discriminator Networks. The X-axis shows the number of iterations, and the Y-axis shows the value of the Loss. The quality of images being generated is improving with each iteration, as is the corresponding loss value. Additionally, the value of the normalized discriminator loss varies similarly for different iterations, showing that the performance of the discriminator is stable throughout the entire process of training, whereas the variations in the values of the generator loss are more, showing that the Generator continuously improves its ability to produce images that can successfully fool the Discriminator.

This shows that the entire process of training is stable, and the GAN is gradually reaching a state where it can produce real-like images of the face.

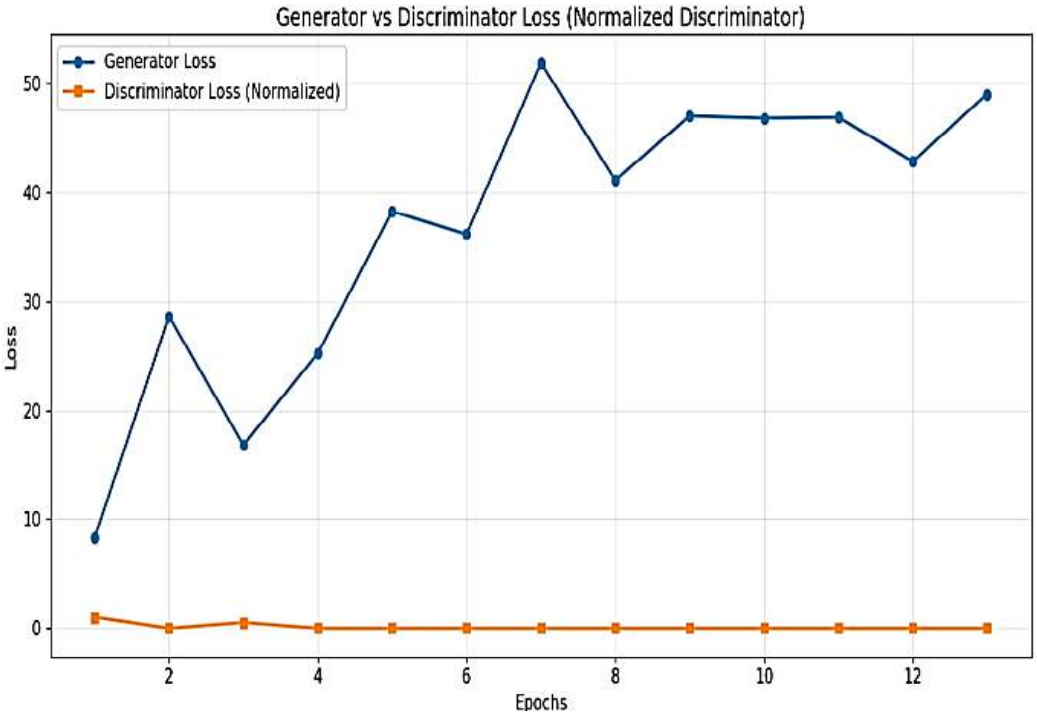

Figure 5: Training loss versus iteration

Table 1: Training Loss of Generator and Across Epochs

Epochs	Generator Loss	Discriminator Loss
1	8.29777431488	0.00024976596785
2	28.58876800	1.01271225855e-07
3	16.80978775	0.0085612472892
4	25.21379852	1.4327762413e-11
5	38.24016952	3.6444751504e-08
6	36.16258239	9.9901673702e-11
7	51.87196731	1.0748372479e-21
8	41.11318206	1.7828529187e-16
9	47.01142501831055	1.170715597e-16
10	46.808837890625	2.41441523e-18

The Generator loss and the Discriminator loss were observed during the training process for a GAN model for the first 10 epochs of training, as shown in Table 1. The generator loss varies erratically for each epoch, first increasing and then constantly decreasing as the model moves towards being able to generate realistic results. Such a process is expected to occur for a GANs as the Generator tries to deceive the Discriminator. The difference here is that the Discriminator loss has a sharp decline followed by a decrease in the nearly zero values approaching 10-7 to 10-18 levels, making the Discriminator very certain in its determination of whether a given sample is real or fake.

The nearly-zero values in the losses indicate that the Discriminator is dominating the training process since it is able to correctly distinguish the fake data as such.

In general, the loss values show that the model is learning because the Generator is getting better and better, and the Discriminator is confident in its results. Such values are normal in the dynamics of a GAN model and typically occur during the early stages of training.

B. Generated New Face Images

The produced output of the trained Generator model is depicted in the following images. It becomes evident that the model has effectively captured the key features of a face, including the positioning of the eyes, symmetry of the face, shape of the hair features, and features of the skin. Even though there are certain small artifacts in the output that suggest a certain blurriness, as well as a few features that might be expected in a moderate-resolution GAN model or a model trained for a few epochs, the overall outlook of the output depicts that the model is indeed capable of replicating a face from a combination of randomly generated features

without the need to trace any specific features from a real-life image or dataset.
Figure 6 shows a single synthetic face generated by the trained GAN model, reflecting that the Generator successfully learns facial structure and texture patterns from the dataset.

- A Single Face Image Generated by the Proposed GAN Model

Figure 6: Single image generated by GAN model

In the below Figure 7 shows a collection of generated face images using different noise inputs, demonstrating output diversity and confirming that the model avoids mode collapse.

- Grid of Various Synthetic Faces Generated by the GAN

Figure 7: Grid of image generated by GAN model

C. Evaluation

The evaluation of the GAN model proposed for this project has been carried out in the context of the analysis of the learning process and the image evaluation for its reality and diversity aspects within the implementation stage of the project.
As a matter of fact, the evaluation of the image for reality and diversity is the most appropriate within the context of the overall objective of the project that aims at image generation with the objective of the generation of realistic human faces.
During the training process, the performance of the model on the loss functions for both Generator & Discriminator was noted. The loss for the Generator varied during the training process. This showed that Generator was learning how to adjust its weights based on how well the generated images looked like reality based on the output of the Discriminator. The loss for the output of the Discriminator decreased significantly with time and approximated a zero-value close to the end of the training process. This showed the confidence of Discriminator in distinguishing real images from a generated sample.
Apart from the analysis of loss, the part of the evaluation of models was the evaluation by visual observation. This was done by generating the images at random points in the epochs of training using random vectors of the latent space. The images were observed for their facial structure, symmetry, clarity, and details of the eyes, nose, and mouth. The improvement of facial coherence of the generated images over a number of epochs was a criterion of how the models had learned.
For diversity, it was necessary that these images be generated using varying amounts of random noise. One of the indicators of diversity in these images would be that they should include variations of facial features, haircut, orientation, and lighting. Variations of these features would also ensure that the

Generator was able to learn an adequately diverse distribution of facial features instead of learning to generate an entire set of repetitive outputs.

In addition, a degree of similarity between such images and actual samples from the training data regarding structure and texture was noted. Although such images contained noticeable artifacts and blurs in certain areas, especially in the background and hairstyle, the composition of the face remained unchanged.

Overall, in this research, a performance evaluation strategy has been adopted where more credence has been assigned to the implementation level assessment than just focusing on numbers. This strategy for analyzing, with the help of observation of loss curves and rigorous analysis, has made it possible to determine the efficacy of learning with the developed GAN for creating human facial images.

VIII. CONCLUSION

This paper presented a Generative Adversarial Network for realistically generating images of human faces. Using a convolution GAN, the presented method is able to learn the distribution of facial information effectively, generating highly realistic images of human faces randomly. The presented method was trained on the Face Mask Lite Dataset, without mask, with the images resized to a resolution of 128 x 128 pixels.

These experimental findings clearly reveal that the Generator is able to produce high-quality facial images with proper structure, symmetry, and variability, and that the Discriminator is playing a proper role in guiding the training procedure through its feedback mechanism in order to form an optimal model that is free from the usual training difficulties associated with GANs despite their nature being generally problematic for training purposes. An essential advantage of this proposed work is its simplicity and efficiency, which can be further beneficial in an academic environment with less computational capability. Although a slight imperfection is present in the synthesized images, their quality is a sign of confirmation that this model performs significantly well for face image generation.

Further, the technical contribution shown in this paper illustrates the comparatively more significant point of instilling ethical consciousness in generated content based on AI. With this emphasis on improvement being expected in face generation models, there will have to be a certain responsibility on part of the development and evaluation process that excludes any potential misuse of the results. In conclusion, this paper has introduced a face generation system based on adversarial learning.

IX. FUTURE SCOPE

Although the proposed model of GAN showed very encouraging results in generating real images of the human face, this research can be extended in several directions. Further research can be done by improving the quality and resolution of the portrayed images, using more advanced models of GAN, including StyleGAN2, StyleGAN3, and StyleGAN-ADA, which allow for higher resolution and more detailed images of faces.

With the present model, it is possible to unconditionally generate faces, but future research may consider generating faces with attributes using conditional GAN methods like StarGAN or AttGAN. In this respect, certain facial attributes such as age, expressions, hairstyle, and facial accessories could be adjusted in order to enhance the usability of this model. Further improvements could thus be derived by including an SR method aimed at optimizing detailed information along with artifacts within generated face images. Furthermore, application of standard quantitative evaluation methods like FID and Inception Score would result in objective model performance evaluation instead of qualitative assessment.

In future studies, more ethical means can be made in which the mechanism can detect the face produced by GAN and not allow them to be used for any malicious activity. Additionally, one can extend this technology to create an application that can allow users to see in real time an image of the face produced by computer-generated faces or any other face generation technology.

CONFLICTS OF INTEREST

The authors declare that they have no conflicts of interest.

REFERENCES

[1] S. Singh, A. Sharma, and R. Gupta, "GAN-based human face generation using CelebA dataset," *International Journal of Computer Applications*, vol. 174, no. 20, pp. 1–6, 2021.

[2] M. Mahiuddin and A. Chowdary, "Synthetic face generation for law-enforcement applications," *IEEE Access*, vol. 8, pp. 123456–123465, 2020.

[3] V. S. K. Katta, P. S. Reddy, and K. R. Rao, "High-resolution face synthesis using DCGAN and ESRGAN," *International Journal of Emerging Technologies in Engineering Research*, vol. 8, no. 6, pp. 45–50, 2020.

[4] Wang, Y. Liu, and J. Zhang, "A comprehensive review of GAN-generated face detection methods," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 4, no. 2, pp. 1–15, 2022. Available from: https://ieeexplore.ieee.org/abstract/document/10604674

[5] R. Reddy and M. Raman, "Conditional GANs for attribute-controlled face generation," *International Journal of Computer Vision and Image Processing*, vol. 10, no. 3, pp. 22–35, 2020.

[6] Goodfellow *et al.*, "Generative adversarial networks," in *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 2672–2680, 2014. Available from: https://tinyurl.com/bdzj9f24

[7] Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional GANs," in *Proc. Int. Conf. Learn. Representations (ICLR)*, 2016. Available from: https://doi.org/10.48550/arXiv.1511.06434

[8] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein GAN," in *Proc. Int. Conf. Mach. Learn. (ICML)*, pp. 214–223, 2017. Available from: https://proceedings.mlr.press/v70/arjovsky17a.html

[9] Gulrajani *et al.*, "Improved training of Wasserstein GANs," in *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 5767–5777, 2017. Available from: https://tinyurl.com/2ppzn37x

[10] T. Miyato *et al.*, "Spectral normalization for generative adversarial networks," in *Proc. Int. Conf. Learn. Representations (ICLR)*, 2018. Available from: https://doi.org/10.48550/arXiv.1802.05957

[11] T. Karras *et al.*, "Progressive growing of GANs for improved quality, stability, and variation," in *Proc. Int. Conf. Learn. Representations (ICLR)*, 2018. Available from: https://doi.org/10.48550/arXiv.1710.10196

[12] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proc. IEEE*

Conf. Comput. Vis. Pattern Recognit. (CVPR), pp. 4401–4410, 2019. Available from: https://tinyurl.com/3tz5p7r9
[13] T. Karras *et al.*, "Analyzing and improving the image quality of StyleGAN," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2020. Available from: https://tinyurl.com/2s4accnp
[14] P. Isola *et al.*, "Image-to-image translation with conditional adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 1125–1134, 2017. Available from: https://tinyurl.com/yupd6858
[15] J. Zhu *et al.*, "Unpaired image-to-image translation using CycleGAN," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, pp. 2223–2232, 2017. Available from: https://tinyurl.com/yc4xzedv
[16] T. Karras *et al.*, "Training generative adversarial networks with limited data," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. Available from: https://tinyurl.com/fuva2t4k
[17] Y. Choi *et al.*, "StarGAN: Unified generative adversarial networks for multi-domain image translation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 8789–8797, 2018. Available from: https://ieeexplore.ieee.org/abstract/document/8545169
[18] T. Karras *et al.*, "Alias-free generative adversarial networks," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. Available from: https://tinyurl.com/yfu6bhdu
[19] Z. He *et al.*, "AttGAN: Facial attribute editing by only changing what you want," *IEEE Transactions on Image Processing*, vol. 28, no. 11, pp. 5464–5478, 2019. Available from: https://ieeexplore.ieee.org/abstract/document/8718508
[20] Härkönen *et al.*, "GANSpace: Discovering interpretable GAN controls," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. Available from: https://tinyurl.com/btchhyt9

ABOUT THE AUTHORS

Supritha P O completed her BE in the year 2017 and M.Tech in the year 2020 from VTU Belagavi. She is an Assistant Professor in the Department of Computer Science & Engineering at SDM Institute of Technology. Area of interest: AI, cloud computing, networking and security.

Prajwal is currently pursuing a B.E. degree in Computer Science and Engineering at SDM Institute of Technology, Ujire, affiliated with Visvesvaraya Technological University (VTU), Belagavi. His areas of interest include deep learning, the Internet of Things (IoT), and Artificial Intelligence (AI)

Saikumar Laxman Pujari. is currently pursuing a B.E. degree in Computer Science and Engineering at SDM Institute of Technology, Ujire, affiliated with Visvesvaraya Technological University (VTU), Belagavi. His areas of interest include deep learning, the Internet of Things (IoT), and Artificial Intelligence (AI)

Siddartha R. is currently pursuing a B.E. degree in Computer Science and Engineering at SDM Institute of Technology, Ujire, affiliated with Visvesvaraya Technological University (VTU), Belagavi. His areas of interest include deep learning, the Internet of Things (IoT), and Artificial Intelligence (AI)

Adarsh Shendage is currently pursuing a B.E. degree in Computer Science and Engineering at SDM Institute of Technology, Ujire, affiliated with Visvesvaraya Technological University (VTU), Belagavi. His areas of interest include deep learning, the Internet of Things (IoT), and Artificial Intelligence (AI)