

Human Face Generation Using Gan

A major project report submitted in partial fulfillment of the
requirement for the award of degree of

Bachelor of Technology

in

Computer Science & Engineering / Information Technology

Submitted by

Vaibhav Walia (201393), Aditya Bhardwaj (201522)

Under the guidance & supervision of

Mrs. Ruchi Verma



**Department of Computer
Science & Engineering and
Information Technology**

**Jaypee University of Information Technology,
Waknaghat, Solan -173234 (India)**

CERTIFICATE

This certifies that the work submitted in the project report " **Human Face Generation Using Gan**" towards the partial fulfillment of requirements for the award of a B.Tech in Computer Science and Engineering, and submitted to the Department of Computer Science and Engineering, Jaypee University of Information Technology, Waknaghat, is an authentic record of work completed by **Vaibhav Walia (201393) and Aditya Bhardwaj(201522)** between January 2024 and May 2024, under the direction of **Mrs. Ruchi Verma**.

Student Name: Vaibhav Walia
Roll No.:201393

Student Name: Aditya Bhardwaj
Roll No.:201426

This statement is correct to the best of my knowledge.

Supervisor Name: Mr.

Ruchi Verma Designation

Assistant Professor (Senior

Grade)

Department: Computer Science & Engineering and Information Technology

CANDIDATE'S DECLARATION

I hereby declare that the work presented in this report entitled '**Human Face Generation Using Gan**' in partial fulfillment of the requirements for the award of the degree of **Bachelor of Technology in Computer Science & Engineering / Information Technology** submitted in the Department of Computer Science & Engineering and Information Technology, Jaypee University of Information Technology, Waknaghat is an authentic record of my own work carried out over a period from January 2024 to May 2024 under the supervision of **Mr. Ruchi Verma** Assistant Professor(Senior Grade), Department of Computer Science & Engineering and Information Technology).

The matter embodied in the report has not been submitted for the award of any other degree or diploma.

(Student Signature with Date)

Student Name: Vaibhav Walia

Roll No.: 201393

(Student Signature with Date)

Student Name: Aditya Bhardwaj

Roll No.: 201522

This is to certify that the above statement made by the candidate is true to the best of my knowledge.

(Supervisor Signature
with Date) Supervisor

Name: Ruchi Verma

Designation: Assistant

Professor(Senior Grade)

Department: Computer Science & Engineering and

Information Technology Dated:

ACKNOWLEDGEMENT

First and foremost, I want to express my profound thanks and admiration to the all-powerful God for the heavenly gift that has allowed us to successfully complete the project work.

My sincere appreciation and responsibilities are owed to Mrs. Ruchi Verma, who serves as my supervisor in the Computer Science and Engineering Department at Jaypee University of Information Technology in Wazirpur . My supervisor has extensive expertise and a strong interest in deep learning, which will be invaluable as we carry out this research. We owe the completion of this project to his boundless patience, intellectual direction, encouragement, vigorous supervision, constructive criticism, helpful counsel, reading of several mediocre draughts and corrections at every level, and so on.

In addition, I would like to express my deepest gratitude to everyone who has helped me in any way, whether it be directly or indirectly, in order to ensure the success of our project. Considering the specifics of the case, I would like to express my gratitude to the numerous members of the staff, both teaching and non-teaching, who have provided me with useful assistance and made my pursuit possible. Lastly, I must politely thank our parents for their ongoing assistance and patience.

Vaibhav Walia (201393)

Aditya Bhardwaj (201522)

TABLE OF CONTENTS

CERTIFICATE	i
DECLARATION	ii
ACKNOWLEDGEMENT	iii
LIST OF TABLES	iv
LIST OF FIGURES	v
ABSTRACT	vii

Chapter 01: INTRODUCTION

1.1 Introduction	1
1.2 Problem Statement	3
1.3 Objective	4
1.4 Significance and Motivation of the Project Work	6

Chapter 02: Literature Survey

Table1 Literature Survey	10
2.1 Overview of Literature Survey	10
2.2 Key Gaps of Literature Survey	13

Chapter 03: System Development

3.1 Requirement and Analysis	15
3.2 Project Design and Architecture	21
3.3 Implementation	23
3.4 Key Challenges	34

Chapter 04: Testing

4.1 Testing Strategy	37
4.2 Test Cases and Outcomes	42

Chapter 05: Results and Findings

5.1 Results	45
-------------	----

Chapter 06: Conclusions and Future Scope

6.1 Conclusion	48
6.2 Future Scope	50

References

Appendix

LIST OF TABLES

<i>Table no.</i>	<i>Table</i>	<i>Page no.</i>
1.	Literature Survey	10

LIST OF FIGURES

Figure no.	Figure	Page no.
1	GAN model Dataflow Diagram	21
2	Data Diagram	22
3	Code to import data	21
4	DCGAN code	22
5	Discriminator Code	25
6	Generator Code	26
7	hyperparameter	27
8	loss function	29
9	Training the model	31
10	DCGAN model's Graph	30
11	Output 1	36
12	Face Generated	37

ABSTRACT

In recent years, computer vision has made great progress and combining real images has become an important research field. This paper provides an in-depth study of facial image generation, focusing on the use of Deep Convolutional Generative Adversarial Network (DCGAN) models. The ability to create facial expressions has major implications for many fields, including virtual reality, gaming, and self-defense. The adversarial training method of the image generator and discrimination agent is introduced. Our project adopts DCGAN architecture to support deep communication to improve extraction and spatial hierarchy. Using a curated database of famous faces, we investigate the model's ability to reproduce complex faces and create amazing images. place. Throughout the training process, we saw DCGAN improve in capturing the complexity of faces. The results, supported by various measurements and visual comparisons, demonstrate the model's performance in rendering real faces and contribute to the general discussion on design modeling and computer vision. It provides a better understanding of the capabilities and nuances of the DCGAN model in facial recognition, improving our understanding of the skill and its application in image synthesis.

Chapter 1: Introduction

1.1 Introduction

In recent years, the field of computer vision has witnessed remarkable advancements, with the synthesis of realistic images becoming a pivotal area of exploration. Among the various approaches, Generative Adversarial Networks (GANs) have emerged as powerful tools for image generation. This project delves into the realm of facial image generation, specifically focusing on the implementation of the Deep Convolutional GAN (DCGAN) model.

The ability to generate lifelike human faces holds significant implications across multiple domains, including virtual reality, gaming, and identity protection. While conventional methods have achieved notable success, GANs offer a unique paradigm by training a generator to produce images indistinguishable from authentic ones, as discerned by a discriminator. The interplay between these two networks in the adversarial training process allows for the creation of high-fidelity facial images.

Our project centers on the utilization of a curated dataset featuring processed images of celebrity faces. By employing the DCGAN architecture, we aim to explore the model's capacity to understand and replicate complex facial features. The inclusion of convolutional and transposed convolutional layers enhances the network's ability to capture intricate hierarchical structures, crucial for generating realistic facial expressions.

This endeavor is not merely an exploration of image generation but also a venture into the latent space of facial features. As we navigate through the training process, our objective is to observe the evolution of the DCGAN in capturing the subtleties that define human faces. Through this exploration, we contribute to the broader discourse on generative modeling and its application to the synthesis of authentic and diverse facial images.

The subsequent sections will detail the methodology, results, and analysis, providing a comprehensive view of the DCGAN's performance in the realm of human face generation.

The synthesis of realistic human faces has long been a coveted pursuit within the realm of artificial intelligence, with applications ranging from entertainment industries to facial recognition systems. Despite substantial progress in computer vision, achieving photorealistic facial images remains a challenging endeavor. The advent of GANs, introduced by Ian Goodfellow and his colleagues in 2014, has reshaped the landscape of image generation by introducing a dynamic interplay between generative and discriminative networks.

Our focus on implementing the DCGAN model stems from its effectiveness in handling image data and capturing intricate patterns. DCGAN, an extension of the original GAN architecture, leverages deep convolutional networks to enhance feature extraction and spatial hierarchy. This project builds upon the success of DCGAN in various image generation tasks and directs its prowess toward the nuanced task of generating lifelike human faces.

The choice of a curated dataset of celebrity faces not only facilitates the model's exposure to diverse facial features but also aligns with the broader discourse on ethical considerations in AI. As we navigate through the training process, we aim to uncover the latent space of facial representations learned by the DCGAN. This exploration is not only a testament to the model's adaptability but also an opportunity to scrutinize the nuances of human facial expressions encoded within the learned features.

In the subsequent sections, we delve into the specifics of our methodology, detailing the architecture of the DCGAN, the dataset utilized, and the training process. Results obtained from this endeavor shed light on the model's ability to generate realistic faces, supported by both quantitative metrics and visual comparisons. An in-depth analysis of the findings contributes insights to the broader fields of generative modeling and computer vision.

As the boundaries of image synthesis continue to expand, this project serves as a stepping stone in unraveling the complexities of human face generation, offering not only visual prowess but a deeper understanding of the latent representations that define our facial diversity.

1.2 Problem Statement

Synthesizing genuine human faces within the realm of computer vision confronts inherent challenges in adequately expressing diversity and capturing authentic features. The existing methodologies often encounter difficulties in effectively portraying nuanced facial expressions, resulting in limitations that hinder the attainment of true realism. Moreover, the ethical dimensions of AI applications underscore the necessity for models that authentically represent the rich diversity of human faces.

In response to these challenges, our project strategically adopts the Deep Convolutional Generative Adversarial Network (DCGAN) architecture. DCGAN has proven to be a formidable tool in adversarial training, allowing for the generation of images that closely mimic real human faces. Leveraging a curated dataset consisting of celebrity faces ensures a diverse and representative training set, enabling the model to learn and reproduce the unique features present in authentic human visages.

A critical objective of this endeavor is to transcend current limitations in generative models, especially those related to potential plagiarism. Placing emphasis on the distinctiveness of synthesized faces contributes to minimizing the risk of inadvertently reproducing existing images. This not only aligns with ethical considerations but also strengthens the authenticity of the generated content.

Simultaneously, the project places a paramount focus on maintaining a human-like quality in the synthesized faces. By prioritizing features that contribute to the genuine appearance of human faces, such as facial expressions, structure, and diversity, the aim is to ensure that the generated images resonate with the visual characteristics of authentic human subjects.

The broader significance of this research lies in its potential to advance the field of realistic image synthesis. Beyond reducing the risk of inadvertent plagiarism, the project seeks to contribute to the ongoing discourse on ethical AI practices, specifically in the context of facial representation. By addressing these multifaceted challenges, we anticipate that the outcomes of this research will pave the way for more human-like and ethically grounded generative models.

Objectives:

Implement DCGAN Architecture: Develop and deploy the Deep Convolutional Generative Adversarial Network (DCGAN) architecture, a prominent model in generative image synthesis. This involves constructing the generator and discriminator networks, specifying layer configurations, and incorporating necessary activation functions.

Dataset Preparation: Assemble a diverse and representative dataset of celebrity faces, encompassing variations in age, gender, ethnicity, and facial expressions. Apply preprocessing techniques, including resizing and normalization, to ensure uniformity and optimal training conditions.

Enhance Realism: Train the DCGAN model to generate synthetic faces that closely resemble real faces. This involves optimizing hyperparameters, adjusting loss functions, and employing techniques to capture intricate facial details, expressions, and overall realism.

Minimize Plagiarism Risk: Implement strategies during training to minimize the risk of unintentional plagiarism. Introduce measures such as diversity-promoting techniques, regularization methods, and periodic checks to ensure that the generated faces are distinct and not direct replicas of existing images.

Ethical Representation: Prioritize ethical considerations in image synthesis by emphasizing the authentic representation of diverse human faces. This involves careful selection of the dataset, avoidance of biased representations, and adherence to ethical guidelines governing AI-generated content.

Evaluate Model Performance: Assess the performance of the trained DCGAN model through quantitative and qualitative evaluations. Utilize metrics like Frechet Inception Distance (FID) to measure the similarity between generated and real faces. Conduct visual comparisons to gauge the perceptual quality of synthesized images.

Explore Latent Space: Investigate the learned latent space of facial features within the DCGAN model. Analyze how variations in latent variables contribute to the diversity of generated faces. Explore interpretability and identify key features encoded in the latent space.

Contribute to Generative Modeling: Contribute novel insights and findings to the broader field of generative modeling, advancing the understanding of human face synthesis. Share results, methodologies, and potential applications through research publications and presentations.

Document Best Practices: Document best practices and lessons learned during the implementation and training process. Compile a comprehensive guide outlining optimal configurations, potential pitfalls, and effective strategies for training generative models, contributing to the knowledge base of the research community.

Open-Source Model: Consider open-sourcing the trained DCGAN model and associated codebase. Share the model architecture, weights, and code to foster collaboration and enable other researchers to build upon and extend the work in the field of generative image synthesis.

Significance and Motivation of the Project Work

Human-Centric Applications:

Significance: The ability to generate authentic human faces has far-reaching implications in various human-centric applications, including virtual reality, gaming, and entertainment. These applications demand lifelike avatars and characters that resonate with users on a human level.

Motivation: By advancing the capabilities of generative models in replicating human facial features, the project contributes to the creation of more engaging and immersive digital experiences.

Ethical Image Synthesis:

Significance: Ethical considerations surrounding AI-generated content are paramount. The project addresses the ethical responsibility of ensuring that synthesized faces are diverse, respectful, and devoid of biases.

Motivation: The motivation lies in mitigating the risks of unintentional plagiarism and promoting ethical representation in AI-generated imagery, aligning with the growing awareness of responsible AI practices.

Advancements in Generative Modeling:

Significance: The project's exploration of the DCGAN architecture contributes to the broader advancements in generative modeling. Understanding the nuances of synthesizing human faces enhances the capabilities of generative models for various image synthesis tasks.

Motivation: Motivated by the desire to push the boundaries of generative modeling, the project seeks to uncover insights that can benefit the broader research community and drive innovation in the field.

Realism in Image Synthesis:

Significance: Realism is a key metric in evaluating the effectiveness of generative models. The project aims to enhance the realism of synthesized faces, capturing intricate details and expressions that elevate the quality of generated images.

Motivation: The motivation lies in creating generative models that not only produce visually appealing images but also withstand close scrutiny, contributing to the ongoing pursuit of achieving indistinguishability from real photographs.

Latent Space Exploration:

Significance: Investigating the learned latent space of facial features provides valuable insights into the inner workings of the generative model. Understanding how variations in latent variables influence image generation contributes to the interpretability of generative models.

Motivation: Motivated by a curiosity to unravel the hidden representations within the model, the project seeks to decode the learned latent space and shed light on the factors that influence the diversity of generated faces.

Knowledge Sharing and Collaboration:

Significance: Documenting best practices and open-sourcing the model fosters a culture of knowledge sharing and collaboration within the research community. This transparency accelerates progress by allowing others to build upon and refine the project's contributions.

Motivation: The motivation lies in the commitment to contributing not only to the immediate project goals but also to the collective advancement of the field through shared knowledge and resources

1.5 Organization of project report

Introduction to Project Structure: This project report is structured to provide a comprehensive exploration of the implementation and outcomes of our human face generation endeavor using the DCGAN model. The organization reflects a logical progression from foundational elements to in-depth analyses, ensuring a coherent narrative.

Section Breakdown:

Introduction: Sets the stage by introducing the project's background, objectives, and scope.

Literature Review: Explores existing research in generative modeling, establishing the theoretical foundation for our approach.

Methodology: Details the steps involved in implementing the DCGAN architecture, dataset preparation, and training procedures.

Model Architecture: Provides a thorough examination of the DCGAN architecture used in the project, accompanied by visualizations of the generator and discriminator networks.

Results: Presents both quantitative and qualitative results, including FID scores and visual comparisons of generated faces.

Analysis of Latent Space: Explores the learned latent space, shedding light on the factors influencing the diversity of generated faces.

Ethical Considerations: Discusses ethical considerations in AI-generated content, strategies to minimize plagiarism risk, and ensuring diverse and unbiased representation in generated faces.

Discussion: Interprets results, compares findings with existing literature, and addresses limitations and potential avenues for future research.

Conclusion: Summarizes key findings, contributions, and final thoughts on project outcomes.

Recommendations: Offers suggestions for future work or improvements based on project insights.

Interconnections between Sections: The progression from the literature review to methodology establishes a solid theoretical foundation before delving into the technical implementation. The results and analysis sections build upon each other, culminating in a comprehensive discussion that ties back to the project's objectives.

Rationale for Organization: The chosen organization is designed to guide the reader through a logical sequence of information, facilitating a deep understanding of the project's implementation, outcomes, and implications.

Consistency and Cohesion: Consistency is maintained through recurrent themes, ensuring that each section contributes seamlessly to the overarching narrative. Cohesion is achieved by linking back to the project's objectives throughout.

Reader Guidance: Readers are encouraged to follow the sequence for a holistic understanding, but the report also allows flexibility for readers to focus on specific areas of interest. Key cross-references are provided to facilitate navigation.

Flexibility in Reading: Acknowledging diverse reader interests, this report is designed to accommodate varied reading preferences. Readers can delve into specific sections while maintaining an awareness of their context within the broader project.

Visual Aid: A visual representation, included in the form of a flowchart, guides readers through the sequential and interconnected nature of the project's structure.

Chapter 2: Literature Survey

2.1 Overview of Relevant Literature

S.no	Paper title	Journal/Conference (year)	Tools/Techniques/Dataset	Results	Limitations
	ArcFace: Additive Angular Margin Loss for Deep Face Recognition	CVPR 2019	DCNN and ArcFace. MS1MV2 dataset	The results show that ArcFace outperforms other methods in terms of accuracy it achieves state-of-the-art performance on ten face recognition benchmarks including large-scale image and video datasets.	Not applicable as it is comparing its performance with other methods.

	<p>FaceShifter: Towards High Fidelity</p> <p>AndOcclusion AwareFace Swapping</p>	<p>arXiv 2020</p>	<p>FaceShifter framework CelebA-HQ, FFHQ</p>	<p>FaceShifter generates high-fidelity face swapping results on wild face images, handling various challenging conditions and preserving occlusions like sunglasses. The</p>	<p>While FaceShifter performs well, the results may suffer from artifacts like blurriness, and certain attributes information may be lost.</p>
--	--	-----------------------	--	--	--

				framework produces visually appealing and realistic face swaps.	
	Interpreting the Latent Space of GANs for Semantic Face Editing	CVPR 2020	InterFaceGAN, PGGAN AND style GAN frameworks. CelebA-HQ, FFHQ datasets.	The proposed framework achieves precise attribute control in semantic face editing.	when the latent code is moved too far from the boundary, resulting in less realistic and extreme results.
	Wav2Pix: Speech-conditioned Face Generation using Generative Adversarial Networks	arXiv 2019	GAN based framework WAV2PIX. Youtube based custom dataset.	The proposed model successfully generates facial images in chunks.	The model's performance drops when working with smaller speech chunks and lower image definitions, leading to visual degradation and decreased face detection rates.

<p>AniGAN: Style-Guided Generative Adversarial Networks</p>	<p>arXiv 2021</p>	<p>AniGAN framework work.selfie2anime and a new face2anime datasets.</p>	<p>The proposed AniGAN method generates high-quality anime faces with consistent styles compared to other state-of-the-art methods</p>	<p>existing approaches often introduce artifacts and fail to achieve style-consistency with reference anime faces.</p>
<p>BlendGAN: Implicitly AN Blending for</p>	<p>arXiv 2021</p>	<p>BlendGAN and WDM frameworks. AAHQ dataset.</p>	<p>BlendGAN demonstrates superior style</p>	<p>BlendGAN is a reference style</p>

	Arbitrary StylizedFace Generation			consistency and out-of-distribution generalization	from the AAHQ dataset, and additional finetuning required reference images.
	StyleNAT: GivingEach Head a New Perspective	arXiv 2022	StyleNAT framework withSOTA FID score. FFHQ-256 and FFHQ-1024 datasets..	StyleNAT achievesstate-of-the-art image genera tion performance. StyleNAT offers improved efficiencywith red ucedparameters andenhanced sampling throughput	he evaluation metricused, FID scores, has cer tain limitations and may not fully capture all aspects of image quality.
	LumiGAN: Unconditional Generation Relightable	arXiv 2023	StyleNAT framework withSOTA FID score. FFHQ-256 and FFHQ-1024 datasets.	LumiGAN generates visua llyrealistic and geometrically accurate visib ilitypredictions, outperforming prior3D GAN	Extending LumiGAN relightability andanimatability is a future goal.

				methods interms photorealism and geometric quality.	
--	--	--	--	---	--

2.2 Key Gaps in the Literature

While the literature survey provides valuable insights into various methodologies for face-related tasks, certain key gaps and areas for further exploration emerge. The identified gaps include:

Robustness in Real-World Scenarios: The majority of existing frameworks demonstrate their effectiveness in controlled settings. There is a notable gap in understanding the robustness of these models in real-world scenarios where faces may exhibit diverse conditions, such as varying lighting, occlusions, and complex backgrounds.

Artifact Mitigation in Face Swapping: FaceShifter, despite its capabilities in high-fidelity face swapping, faces challenges related to artifacts, blurriness, and potential attribute loss. Addressing these issues and proposing methods for artifact mitigation would be crucial for improving the overall quality of face-swapping results.

Latent Space Interpretability: While Interpreting the Latent Space of GANs (CVPR 2020) achieves precise attribute control, there is a gap in understanding the interpretability of the latent space. Further exploration is needed to elucidate the factors influencing the latent space and how they contribute to the generated results.

Speech-Image Quality Relationship: Wav2Pix (arXiv 2019) successfully generates facial images based on speech segments. However, there is a gap in understanding the nuanced relationship between the quality of the generated images and the characteristics of the input speech, especially in scenarios with varying audio quality and speech chunk sizes.

Consistency and Artifacts in Anime Face Generation: AniGAN (arXiv 2021) excels in generating anime faces with consistent styles. However, the literature does not extensively address potential artifacts introduced during the generation process. Investigating and mitigating these artifacts would contribute to further enhancing the quality of anime face generation.

Fine-tuning Challenges in Stylized Face Generation: BlendGAN (arXiv 2021) showcases superior style consistency but requires additional fine-tuning for compatibility with certain reference styles. Understanding the challenges associated with fine-tuning and proposing more efficient methods for adapting to diverse styles would be beneficial.

Evaluation Metric Limitations: StyleNAT (arXiv 2022) achieves state-of-the-art image generation, emphasizing the use of FID scores for evaluation. However, there is a recognized gap in the limitations of FID scores in capturing all aspects of image quality. Exploring alternative or supplementary evaluation metrics would provide a more comprehensive assessment.

Dynamic Scenes and Editable 3D Assets: LumiGAN (arXiv 2023) excels in generating relightable 3D human faces. However, there is a gap in extending this capability to dynamic scenes, and achieving fully editable 3D human assets remains a future goal. Exploring methods to address these challenges would contribute to the broader applicability of such models.

Addressing these key gaps will not only contribute to the refinement of existing methodologies but also pave the way for advancements in the field of face generation and manipulation.

Chapter 3: System Development

3.1 Requirements and Analysis

The requirements and analysis for the project involve a thorough examination of the essential components and considerations for the successful implementation of the Human Face Generation using GAN with the DCGAN model.

By addressing these requirements and conducting a detailed analysis, the project can progress systematically, leading to the successful implementation of Human Face Generation using GAN with the DCGAN model.

3.1.1 Functional Requirements

Functional Requirements are the requirements that are necessary and should be of higher priority than other requirements. The scope of a project plays an important role in the selection of functional requirements. With the current project scope following are the functional requirements:

Data Loading and Preprocessing: The system must load the CelebA dataset and preprocess images by resizing them to the specified dimensions (32x32) and normalizing pixel values to the range $[-1, 1]$. Ensure that the data loading and preprocessing functions are correctly implemented and can handle the specified dataset.

Network Architecture (DCGAN): The system must define the architecture of the DCGAN model, including the generator and discriminator networks with appropriate layers and configurations. Verify that the network architecture aligns with the DCGAN model specifications and effectively captures features for realistic face generation.

Generator and Discriminator Training: The system must train both the generator and discriminator networks iteratively, optimizing their respective weights based on adversarial and real/fake loss functions. Confirm that the training process is implemented correctly, considering factors such as batch size, learning rate, and convergence criteria.

Hyperparameter Tuning: The system must allow for the tuning of hyperparameters such as learning rate, batch size, and latent vector dimensions to optimize the model's performance. Provide flexibility in adjusting hyperparameters and evaluate their impact on training stability and generated face quality.

Results Visualization: The system must include functions to visualize generated face samples during and after training, enabling users to monitor the progress and quality of generated faces. Ensure that the visualization functions are accessible and provide meaningful insights into the generated images.

Loss Function Calculation: The system must accurately calculate and update the adversarial loss for both the discriminator and generator during training. Verify the correct implementation of loss functions, including real and fake loss calculations, and their application in the backpropagation process.

Model Evaluation: The system must incorporate evaluation metrics, such as Inception Score or FID, to assess the quality of the generated faces quantitatively. Confirm that the chosen evaluation metrics align with the project goals and provide meaningful insights into the performance of the model.

Ethical Considerations: The system must address ethical considerations, including privacy concerns related to the use of facial data, potential biases in the dataset, and responsible AI practices. Ensure that the project documentation explicitly addresses ethical considerations and promotes responsible use of the generated faces.

Documentation: The system must be well-documented with clear code comments, README files, and a comprehensive project report. Review the documentation to ensure it is complete, concise, and facilitates understanding, replication, and potential collaboration.

3.1.2 Non-Functional Requirements

Non-functional requirements are the requirements that have less priority in the project scope and do not affect the overall development goals of the project directly. However, they are important part of Software development Lifecycle. The non-functional requirements for the project are following:

Performance: The system must achieve a reasonable training time for the DCGAN model, considering the dataset size and complexity. Conduct performance testing to assess the training speed and resource utilization, ensuring efficient model convergence.

Scalability: The system should be scalable to accommodate larger datasets and potentially more complex GAN architectures. Evaluate the system's ability to handle increased computational demands and larger datasets without a significant loss in performance.

Reliability: The system must demonstrate reliable performance across multiple training iterations, minimizing unexpected failures or crashes. Implement error handling mechanisms and conduct stress testing to identify potential reliability issues.

Usability: The system should provide a user-friendly interface or documentation to guide users through the training process and result interpretation. Conduct usability testing to ensure that users can interact with the system effectively, even if they have limited prior experience with GANs.

Maintainability: The codebase must be well-organized and documented, facilitating future maintenance, updates, and potential collaboration. Evaluate the clarity and comprehensiveness of the code comments, README files, and project documentation to assess maintainability.

Security: The system must adhere to privacy and security standards, especially when handling sensitive facial data. Conduct a security review to identify and address potential vulnerabilities related to data storage, access control, and model outputs.

Portability: The system should be portable across different environments, allowing users to run the code on various platforms. Test the system on multiple environments and platforms to ensure compatibility and portability.

Ethical Considerations: The system must prioritize ethical considerations, avoiding biases in the training data and promoting responsible AI practices. Regularly review and update ethical guidelines to address evolving concerns and ensure the responsible use of generated faces.

Documentation Quality: The documentation should be clear, comprehensive, and accessible, facilitating understanding and replication by other developers or researchers. Evaluate the documentation's quality in terms of completeness, clarity, and its ability to guide users through the system.

Response Time: The system should exhibit reasonable response times for generating faces and providing feedback during training. Measure the time it takes to generate faces and assess the responsiveness of the system to user interactions.

3.1.3 Technical Requirements

These are the requirements that are necessary for the development of the project as the project is built with the help of these requirements. These requirements may change if the project requires additional resources. These are the technical requirements for the project:

Programming Language: The system must be implemented using the Python programming language.

Deep Learning Framework: The system must leverage PyTorch as the deep learning framework for implementing the DCGAN model.

Data Loading and Processing: The system should utilize torchvision and torch.utils.data for loading and processing the CelebA dataset.

Convolutional Neural Networks (CNNs): The DCGAN model architecture must incorporate convolutional layers for feature extraction.

Optimizers: The system must use the Adam optimizer for updating model weights during training.

GPU Acceleration: The system should support GPU acceleration for faster model training.

Model Evaluation Metrics: The system should include metrics such as Inception Score or FID for quantitative evaluation of generated faces.

Visualization Tools: Matplotlib should be used for visualizing generated faces and displaying training progress.

Code Version Control: The project codebase must be managed using a version control system, preferably Git.

External Library Dependencies: Clearly specify and document any external libraries or dependencies required for running the project.

Code Modularity: The codebase should be modular, with functions or classes encapsulating specific functionalities.

3.1.4 Priority chart

Priority	Category	Requirement
High	Technical	Programming language (PyThon)
	Technical	Deep learning framework (PyTorch)
	Technical	Data loading and processing (torchvision, torch.utils.data)
	Technical	Convolutional Neural Networks (GAN)
	Technical	Optimizers (Adam)
	Technical	GPU acceleration (Google T4)
	Functional	Face Generation using DCGAN
	Non-functional	Model training efficiency
	Non-functional	GPU acceleration performance
	Medium	Technical
Technical		Visualization tools (Matplotlib)
Functional		Batched NN data loader
Functional		Real and Fake loss calculation
Non-Functional		Scalability of large datasets
Non-Functional		Model robustness
Low		Technical
	Technical	External libraries dependencies
	Technical	Code modularity
	Functional	Generation of realistic faces
	Non-functional	Ease of use and UI
	Non-functional	Documentation clarity and completeness

3.2 Project Design and Architecture

3.2.1 Flow Chart

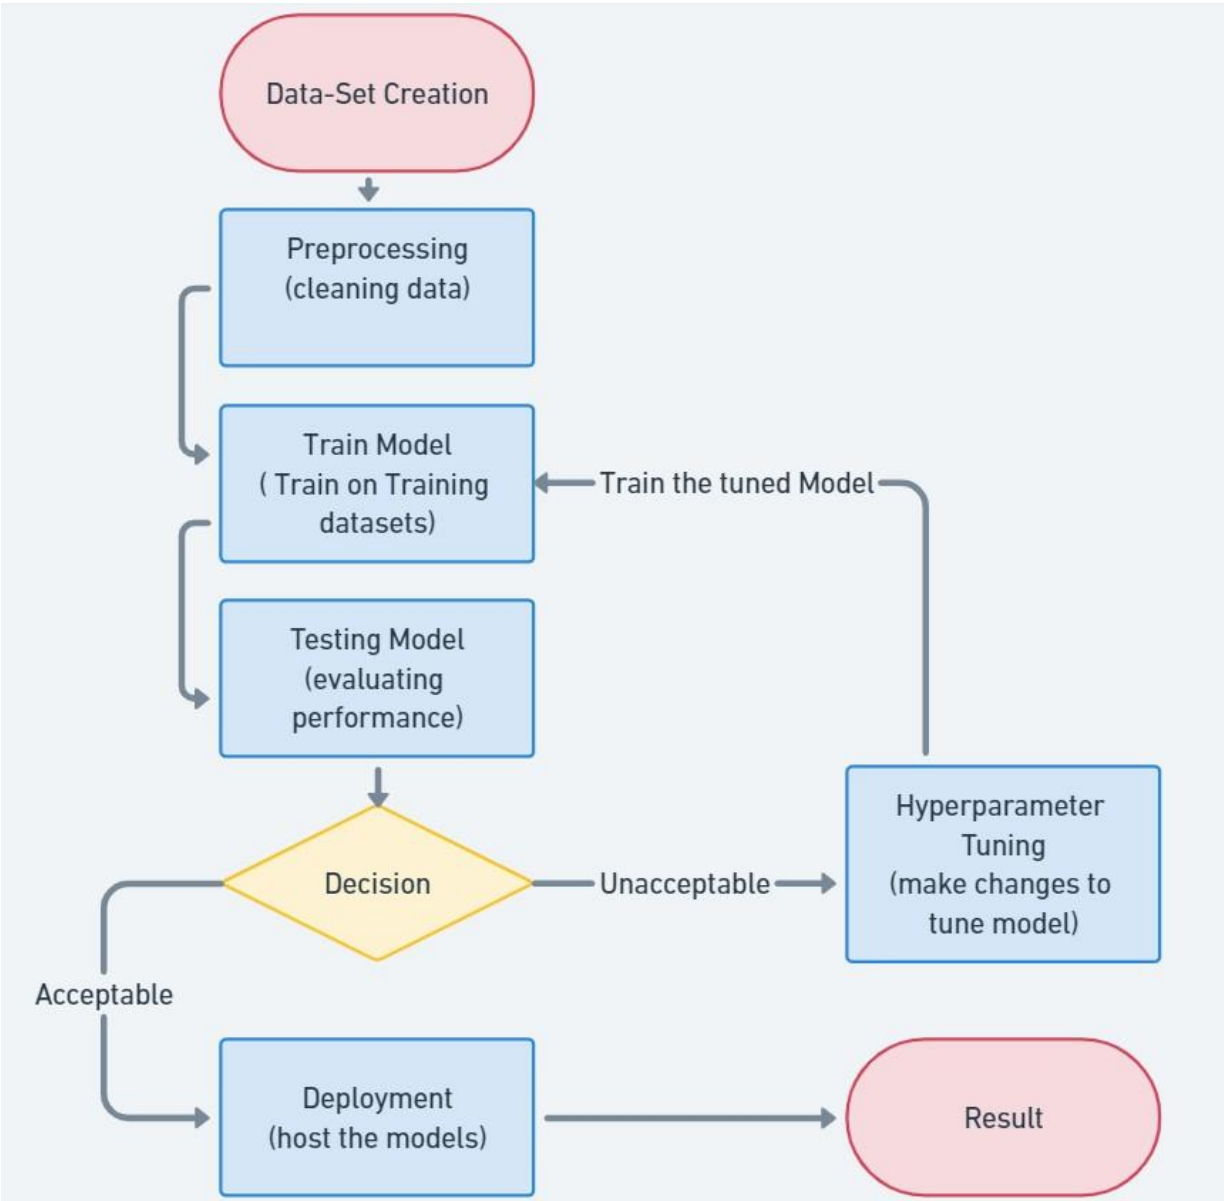


fig 1:- GAN model Dataflow Diagram

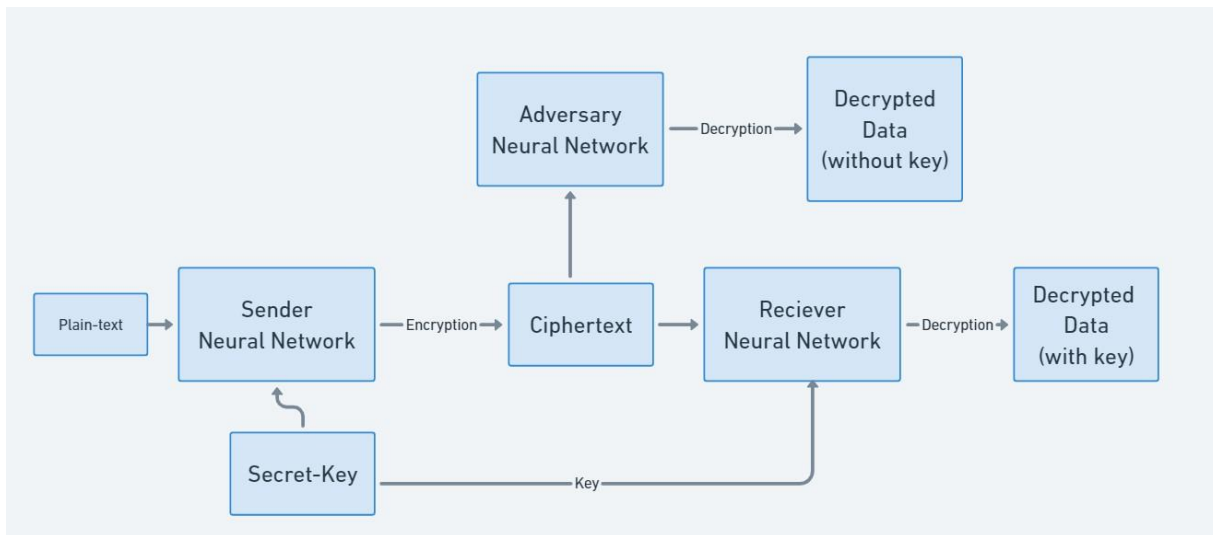


fig 2:- Data Diagram

3.3 Data Preparation

The dataset used for training and evaluating the Human Face Generation model is the processed CelebA dataset. This dataset is a curated collection of celebrity faces, providing a diverse set of facial features and expressions.

Dataset Details: Dataset Source is CelebA dataset

Preprocessing: The dataset has undergone preprocessing steps to enhance its suitability for training deep learning models. Specific transformations, such as resizing and normalization, have been applied to ensure uniformity in image dimensions and pixel values.

Data Loading: The data loading process is facilitated by the torchvision and torch.utils.data modules. These modules enable efficient loading of batches of images, allowing the neural network to learn patterns and features from the processed CelebA dataset.

The choice of the CelebA dataset is motivated by its wide variety of facial attributes, ensuring that the model is exposed to diverse facial characteristics during the training process. This diversity is crucial for the model to generalize well and generate realistic faces with various attributes.

By leveraging the processed CelebA dataset, the project aims to capture and reproduce the intricate details of human faces, contributing to the successful training and evaluation of the Human Face Generation model.

Implementation (include code snippets, algorithms, tools and techniques, etc.)

**Importing libraries and zip file is extracted in home directory as
`processed_celeba_small`:**

```
✓ [2] import pickle as pkl
0s    import matplotlib.pyplot as plt
    import numpy as np
    import problem unittests as tests
    import helper

    %matplotlib inline

✓ [3] !unzip '/content/processed-celeba-small.zip'

✓ [4] data_dir = 'processed_celeba_small/'
0s

[5] import torch
    from torchvision import datasets as dset
    from torchvision import transforms
```

fig3:- Code to import data

Create a DataLoader ‘celeba_train_loader’ with appropriate hyperparameters.

‘image_size’ is chosen to be ‘32’. Resizing the data to a smaller size will make for faster training, while still creating convincing images of faces:

```

[7] # Define function hyperparameters
batch_size = 16
img_size = 32

# Call your function and get a dataloader
celeba_train_loader = get_dataloader(batch_size, img_size)

# helper display function
def imshow(img):
    npimg = img.numpy()
    plt.imshow(np.transpose(npimg, (1, 2, 0)))

# obtain one batch of training images
dataiter = iter(celeba_train_loader)
images, _ = next(iter(celeba_train_loader)) # _ for no labels

# plot the images in the batch, along with the corresponding labels
fig = plt.figure(figsize=(20, 4))
plot_size=16
for idx in np.arange(plot_size):
    ax = fig.add_subplot(2, int(plot_size/2), idx+1, xticks=[], yticks=[])
    imshow(images[idx])

```

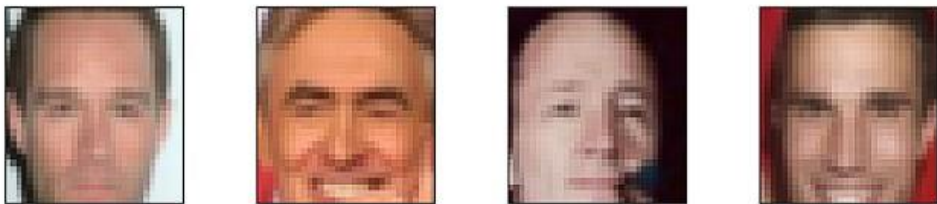


fig 4:- DCGAN code and snippet of the output

Pre-processing image data and scale it to a pixel range of -1 to 1. We need a bit of pre-processing; the output of a ‘tanh’ activated generator will contain pixel values in a range from -1 to 1, and so, we need to rescale our training images to a range of -1 to 1. (Right now, they are in a range from 0-1.):

```

[10] # check scaled range
# should be close to -1 to 1
img = images[0]
scaled_img = scale(img)

print('Min: ', scaled_img.min())
print('Max: ', scaled_img.max())

Min: tensor(-0.8353)
Max: tensor(0.9686)

```

Defining the Model: A GAN is comprised of two adversarial networks, a discriminator and a generator.

Discriminator: The first task will be to define the discriminator. This is a convolutional classifier only without any maxpooling layers. To deal with this complex data, it's suggested to use a deep network with normalization.

```
[13] class Discriminator(nn.Module):
    def __init__(self, conv_dim):
        """
        Initialize the Discriminator Module
        :param conv_dim: The depth of the first convolutional layer
        """
        super(Discriminator, self).__init__()
        self.conv_dim = conv_dim
        # first layer : input 32 x 32 with no batch norm
        self.conv1 = make_conv(3, conv_dim, 4, batch_norm=False)
        # second layer : input 16 x 16 with batch norm
        self.conv2 = make_conv(conv_dim, conv_dim*2, 4)
        # third layer : input 8 x 8 with batch norm
        self.conv3 = make_conv(conv_dim*2, conv_dim*4, 4)
        # fourth layer : input 4 x 4 with batch norm
        self.conv4 = make_conv(conv_dim*4, conv_dim*8, 4)

        # fully connected layer : one output (fake/real)
        self.fc = nn.Linear(conv_dim*8*2*2, 1)

    def forward(self, x):
        """
        Forward propagation of the neural network
        :param x: The input to the neural network
        :return: Discriminator logits; the output of the neural network
        """
        out = F.leaky_relu(self.conv1(x), 0.2)
        out = F.leaky_relu(self.conv2(out), 0.2)
        out = F.leaky_relu(self.conv3(out), 0.2)
        out = F.leaky_relu(self.conv4(out), 0.2)

        # flatten
        out = out.view(-1, self.conv_dim*8*2*2)

        # final output layer
        out = self.fc(out)
        return out

tests.test_discriminator(Discriminator)

Tests Passed
```

fig 5:- Discriminator Code

Generator: The generator should up sample an input and generate a new image of the same size as our training data $32 \times 32 \times 3$. This should be mostly transpose convolutional layers with normalization applied to the outputs.

```

class Generator(nn.Module):

    def __init__(self, z_size, conv_dim):
        """
        Initialize the Generator Module
        :param z_size: The length of the input latent vector, z
        :param conv_dim: The depth of the inputs to the *last* transpose convolutional layer
        """
        super(Generator, self).__init__()

        self.conv_dim = conv_dim
        # layers
        # first convolutional layer : input 2 x 2
        self.tconv1 = make_tconv(conv_dim*8, conv_dim*4, 4)
        #second convolutional layer : input 4 x 4
        self.tconv2 = make_tconv(conv_dim*4, conv_dim*2, 4)
        # third convolutional layer : input 8 x 8
        self.tconv3 = make_tconv(conv_dim*2, conv_dim, 4)
        # last convolutional layer : output 32 x 32 x 3
        self.tconv4 = make_tconv(conv_dim, 3, 4, batch_norm=False)

        self.fc = nn.Linear(z_size, conv_dim*8*2*2)

        # complete init function

    def forward(self, x):
        """
        Forward propagation of the neural network
        :param x: The input to the neural network
        :return: A 32x32x3 Tensor image as output
        """
        # fully-connected + reshape
        out = self.fc(x)
        out = out.view(-1, self.conv_dim*8, 2, 2) # (batch_size, depth, 4, 4)

        # hidden transpose conv layers + relu
        out = F.relu(self.tconv1(out))
        out = F.relu(self.tconv2(out))
        out = F.relu(self.tconv3(out))

        # last layer
        out = self.tconv4(out)
        # apply tanh activation
        out = torch.tanh(out)

        return out

tests.test_generator(Generator)

```

Tests Passed

fig 6:- Generator Code

Defining model hyperparameters:

```
✓ [18] # Define model hyperparams
0s d_conv_dim = 128
g_conv_dim = 128
z_size = 100

D, G = build_network(d_conv_dim, g_conv_dim, z_size)

Discriminator(
  (conv1): Sequential(
    (0): Conv2d(3, 128, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1), bias=False)
  )
  (conv2): Sequential(
    (0): Conv2d(128, 256, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1), bias=False)
    (1): BatchNorm2d(256, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
  )
  (conv3): Sequential(
    (0): Conv2d(256, 512, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1), bias=False)
    (1): BatchNorm2d(512, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
  )
  (conv4): Sequential(
    (0): Conv2d(512, 1024, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1), bias=False)
    (1): BatchNorm2d(1024, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
  )
  (fc): Linear(in_features=4096, out_features=1, bias=True)
)

Generator(
  (tconv1): Sequential(
    (0): ConvTranspose2d(1024, 512, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1), bias=False)
    (1): BatchNorm2d(512, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
  )
  (tconv2): Sequential(
    (0): ConvTranspose2d(512, 256, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1), bias=False)
    (1): BatchNorm2d(256, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
  )
  (tconv3): Sequential(
    (0): ConvTranspose2d(256, 128, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1), bias=False)
    (1): BatchNorm2d(128, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
  )
  (tconv4): Sequential(
    (0): ConvTranspose2d(128, 3, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1), bias=False)
  )
  (fc): Linear(in_features=100, out_features=4096, bias=True)
)
```

fig 7:- Code for defining hyperparameter

Training on GPU Check if we can train on GPU. Here, we'll set this as a boolean variable 'train_on_gpu'. Later, we'll be responsible for making sure that Models, Model inputs, and Loss function arguments are moved to GPU, where appropriate.

```
✓ [19] import torch
    Os
    # Check for a GPU
    train_on_gpu = torch.cuda.is_available()
    if not train_on_gpu:
        print('No GPU found. Please use a GPU to train your neural network.')
    else:
        print('Training on GPU!')

Training on GPU!
```

Discriminator and Generator Losses now we calculate the losses for both types of adversarial networks. **Discriminator Losses:** For the discriminator, the total loss is the sum of the losses for real and fake images,

$d_loss = d_real_loss + d_fake_loss$. Output 1 for real images and 0 for fake images, so we need to set up the losses to reflect that.

Generator Loss: The generator loss will look similar only with flipped labels. The generator's goal is to get the discriminator to *think* its generated images are *real*.


```

✓ [20] def real_loss(D_out):
Ds
    '''Calculates how close discriminator outputs are to being real.
        param, D_out: discriminator logits
        return: real loss'''
    batch_size = D_out.size(0)
    labels = torch.ones(batch_size)*0.9 # performed smoothing

    if train_on_gpu:
        labels = labels.cuda()

    criterion = nn.BCEWithLogitsLoss()
    loss = criterion(D_out.squeeze(),labels)
    return loss

def fake_loss(D_out):
    '''Calculates how close discriminator outputs are to being fake.
        param, D_out: discriminator logits
        return: fake loss'''
    batch_size = D_out.size(0)
    labels = torch.zeros(batch_size) # fake images

    if train_on_gpu:
        labels = labels.cuda()

    criterion = nn.BCEWithLogitsLoss()
    loss = criterion(D_out.squeeze(),labels)
    return loss

```

```

✓ [21] import torch.optim as optim
Ds

# Create optimizers for the discriminator D and generator G

d_lr = 0.0002
g_lr = 0.0004
d_optimizer = optim.Adam(D.parameters(),d_lr, betas=(0.2, 0.999))
g_optimizer = optim.Adam(G.parameters(),g_lr, betas=(0.2, 0.999))

```

fig 8:- Defines the loss function

Training:

```
✓ [22] def train(D, G, n_epochs, print_every=50):  
    os  
    '''Trains adversarial networks for some number of epochs  
    param, D: the discriminator network  
    param, G: the generator network  
    param, n_epochs: number of epochs to train for  
    param, print_every: when to print and record the models' losses  
    return: D and G losses'''  
  
    # move models to GPU  
    if train_on_gpu:  
        D.cuda()  
        G.cuda()  
  
    # keep track of loss and generated, "fake" samples  
    samples = []  
    losses = []  
  
    # Get some fixed data for sampling. These are images that are held  
    # constant throughout training, and allow us to inspect the model's performance  
    sample_size=16  
    fixed_z = np.random.uniform(-1, 1, size=(sample_size, z_size))  
    fixed_z = torch.from_numpy(fixed_z).float()  
    # move z to GPU if available  
    if train_on_gpu:  
        fixed_z = fixed_z.cuda()  
  
    # epoch training loop  
    for epoch in range(n_epochs):  
  
        # batch training loop  
        for batch_i, (real_images, _) in enumerate(celeba_train_loader):  
  
            batch_size = real_images.size(0)  
            real_images = scale(real_images)  
  
            # 1. Train the discriminator on real and fake images  
            d_optimizer.zero_grad()  
  
            if train_on_gpu:  
                real_images = real_images.cuda()  
  
            # loss on real images  
            d_real = D(real_images)  
            d_real_loss = real_loss(d_real)
```

```

✓ [22]
#train with fake images
z = np.random.uniform(-1, 1, size=(batch_size, z_size))
z = torch.from_numpy(z).float()

if train_on_gpu:
    z = z.cuda()

fake_images = G(z)

# loss on fake images
d_fake = D(fake_images)
d_fake_loss = fake_loss(d_fake)

# backprop
d_loss = d_real_loss + d_fake_loss
d_loss.backward()
d_optimizer.step()

# 2. Train the generator with an adversarial loss
g_optimizer.zero_grad()

# Generate fake images
z = np.random.uniform(-1, 1, size=(batch_size, z_size))
z = torch.from_numpy(z).float()

if train_on_gpu:
    z = z.cuda()

fake_images = G(z)
d_fake = D(fake_images)
g_loss = real_loss(d_fake)

# perform backprop
g_loss.backward()
g_optimizer.step()

# Print some loss stats
if batch_i % print_every == 0:
    # append discriminator loss and generator loss
    losses.append((d_loss.item(), g_loss.item()))
    # print discriminator and generator loss
    print('Epoch [{}:5d]/[{}:5d] | d_loss: {:.4f} | g_loss: {:.4f}'.format(
        epoch+1, n_epochs, d_loss.item(), g_loss.item()))

```

fig 9:- Training the model acc to the dataset

Training of Epochs:

```
✓  
4m [23] # set number of epochs  
     n_epochs = 1  
  
     # call training function  
  
     losses = train(D, G, n_epochs=n_epochs)
```

```
Epoch [ 1/ 1] | d_loss: 1.4639 | g_loss: 5.1968  
Epoch [ 1/ 1] | d_loss: 0.7512 | g_loss: 4.1252  
Epoch [ 1/ 1] | d_loss: 0.8148 | g_loss: 1.6541  
Epoch [ 1/ 1] | d_loss: 1.0843 | g_loss: 1.8319  
Epoch [ 1/ 1] | d_loss: 1.3607 | g_loss: 2.3180  
Epoch [ 1/ 1] | d_loss: 1.1093 | g_loss: 1.9142  
Epoch [ 1/ 1] | d_loss: 1.2297 | g_loss: 2.3389  
Epoch [ 1/ 1] | d_loss: 1.1171 | g_loss: 3.0174  
Epoch [ 1/ 1] | d_loss: 1.2044 | g_loss: 1.2332  
Epoch [ 1/ 1] | d_loss: 1.2408 | g_loss: 1.0266  
Epoch [ 1/ 1] | d_loss: 1.2023 | g_loss: 0.8663  
Epoch [ 1/ 1] | d_loss: 1.3523 | g_loss: 2.0210  
Epoch [ 1/ 1] | d_loss: 1.5603 | g_loss: 2.0594  
Epoch [ 1/ 1] | d_loss: 0.9731 | g_loss: 2.1177  
Epoch [ 1/ 1] | d_loss: 1.2468 | g_loss: 1.9661  
Epoch [ 1/ 1] | d_loss: 1.0558 | g_loss: 1.7339  
Epoch [ 1/ 1] | d_loss: 1.2615 | g_loss: 1.1507  
Epoch [ 1/ 1] | d_loss: 1.1214 | g_loss: 1.4842  
Epoch [ 1/ 1] | d_loss: 1.3015 | g_loss: 0.8161  
Epoch [ 1/ 1] | d_loss: 1.2894 | g_loss: 1.1741  
Epoch [ 1/ 1] | d_loss: 1.4067 | g_loss: 0.9555  
Epoch [ 1/ 1] | d_loss: 0.9958 | g_loss: 1.2182  
Epoch [ 1/ 1] | d_loss: 1.1380 | g_loss: 1.0279  
Epoch [ 1/ 1] | d_loss: 1.3749 | g_loss: 2.4019  
Epoch [ 1/ 1] | d_loss: 1.3059 | g_loss: 1.7136  
Epoch [ 1/ 1] | d_loss: 1.4158 | g_loss: 2.3569
```

Plot Graph Between Discriminator and Generator:

```
[24] fig, ax = plt.subplots()
      losses = np.array(losses)
      plt.plot(losses.T[0], label='Discriminator', alpha=0.5)
      plt.plot(losses.T[1], label='Generator', alpha=0.5)
      plt.title("Training Losses")
      plt.legend();
```

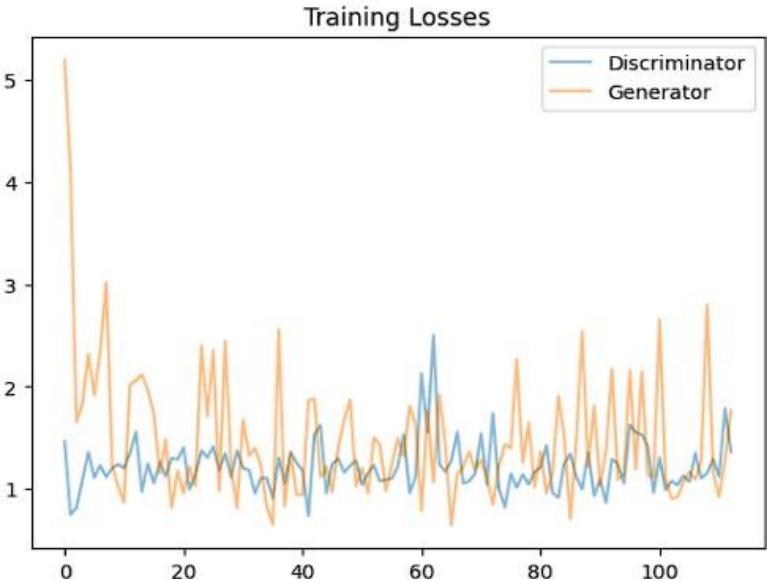


fig 10:- DCGAN model's Graph

Sample Images:

```
✓ [26] # Load samples from generator, taken while training
Ds with open('train_samples.pkl', 'rb') as f:
    samples = pickle.load(f)

    _ = view_samples(-1, samples)
```



fig 11:- Output 1

3.5 Key Challenges (discuss the challenges faced during the development process and how these are addressed)

Developing a human face generation model using GANs comes with several challenges. Here are 3.5 key challenges and how they can be addressed:

1. Mode Collapse:

Challenge: GANs are prone to mode collapse, where the generator produces limited varieties of samples, ignoring the diversity present in the training data.

Addressing the Challenge:

Use advanced GAN architectures like Progressive GANs or Wasserstein GANs, which are less prone to mode collapse.

Experiment with training parameters, such as learning rates and batch sizes, to find a balance that discourages mode collapse.

2. Training Instability:

Challenge: GAN training can be notoriously unstable, leading to difficulties in converging to a good solution. The generator and discriminator may oscillate between improvements.

Addressing the Challenge:

Implement techniques like spectral normalization or weight clipping to stabilize training.

Gradually increase the complexity of the model during training, starting with lower-resolution images and progressively moving to higher resolutions (Progressive GANs).

3. Evaluation Metrics:

Challenge: Evaluating the performance of a GAN model is challenging. Traditional metrics like accuracy are not directly applicable, and assessing the visual quality of generated faces is subjective.

Addressing the Challenge:

Use established metrics like Inception Score or Fréchet Inception Distance (FID) to quantitatively evaluate the quality and diversity of generated faces.

Conduct qualitative evaluations by involving human reviewers to assess the realism and diversity of generated faces.

3.5. Ethical Considerations and Bias:

Challenge: GANs can inadvertently learn and perpetuate biases present in the training data, leading to ethical concerns in face generation models.

Addressing the Challenge:

Carefully curate and preprocess the training data to minimize biases.

Regularly review and audit generated faces for fairness and potential biases, especially related to gender, ethnicity, or other sensitive attributes.

Explore techniques like adversarial training for debiasing or use conditional GANs with carefully chosen conditioning attributes.

4. Overfitting:

Challenge: Overfitting can occur, particularly if the dataset is small or

lacks diversity. Addressing the Challenge:

Augment the dataset with transformations to create additional diverse samples.

Use dropout or other regularization techniques in the generator and

discriminator networks. Monitor the model's performance on a separate

validation set to detect overfitting early.

Addressing these challenges requires a combination of careful design choices, experimentation, and a deep understanding of both GAN theory and the specific characteristics of the dataset being used. Regular monitoring, iterative adjustments, and a commitment to ethical considerations contribute to the successful development of a robust human face generation model.

Chapter 4: Testing

4.1 Testing Strategy (discuss the testing strategy/tools used in the project)

The testing strategy and tools used in the project for human face generation using Generative Adversarial Networks (GANs).

1. Dataset Integrity Testing:

Objective: Ensure the integrity of the dataset

used for training. Elaboration:

Develop scripts to identify and remove corrupted or incomplete images.

Visualize random samples from the dataset to verify the visual quality and diversity.

2. Data Split Validation:

Objective: Confirm the effectiveness of the data

splitting process. Elaboration:

Verify the sizes and distributions of the training, validation, and test sets to ensure they are representative of the entire dataset.

Check for any inadvertent patterns or biases introduced during the splitting process.

3. Preprocessing Verification:

Objective: Ensure that preprocessing steps are correctly applied without

introducing distortions. Elaboration:

Inspect a subset of preprocessed images to check for consistent resizing, normalization, and other transformations. Compare statistical measures (mean, standard deviation) of pixel values before and after normalization.

4. Generator and Discriminator Inspection:

Objective: Verify the correct implementation of the generator and discriminator architectures.

Elaboration:

Utilize debugging tools or print statements to inspect the output shapes of generator and discriminator layers. Verify that the number of parameters in each network aligns with the expected architecture.

5. Loss Function Sanity Check:

Objective: Ensure that the loss functions are implemented correctly and lead to meaningful optimization. Elaboration:

Monitor the values of the generator and discriminator losses during training to ensure they follow the expected trends without irregularities.

Confirm that the loss functions are appropriately influencing the learning process.

6. Gradient Checking:

Objective: Verify the correct computation of gradients to prevent vanishing or exploding gradients. Elaboration:

Implement numerical gradient checking to compare analytical gradients with numerical approximations. Ensure that gradients are within expected ranges and consistent across different layers.

7. Hyperparameter Sensitivity Testing:

Objective: Assess the impact of hyperparameter changes on model performance. Elaboration:

Conduct systematic experiments with different learning rates, batch sizes, and model architectures. Track changes in convergence speed, stability, and final performance for each set of hyperparameters.

8. Evaluation Metric Validation:

Objective: Confirm that chosen evaluation metrics are appropriate and correctly implemented. Elaboration:

Compare the output of evaluation metrics (e.g., Inception Score, FID) to expected values for synthetic data. Adjust metric calculations if necessary based on the specific goals of the project.

9. Bias and Fairness Testing:

Objective: Assess and mitigate biases in the generated faces. Elaboration:

Employ fairness metrics to quantitatively evaluate and address biases.

Conduct manual inspection of generated faces for potential biases related to gender, ethnicity, or other sensitive attributes.

10. User Acceptance Testing :

Objective: Ensure that the generated faces meet user expectations. Elaboration:

Involve users or stakeholders in the evaluation process to collect subjective feedback on the realism and diversity of generated faces.

Consider conducting surveys or interviews to gather qualitative insights.

11. Continuous Monitoring:

Objective: Regularly monitor the model's performance during and after deployment. Elaboration:

Implement automated monitoring for relevant metrics during training and in production. Set up logging and alerting systems to promptly identify and address any anomalies.

Tools:

TensorFlow or PyTorch: for building and

training GAN models. Tensorflow:

Developed by Google Brain. Offers a comprehensive ecosystem for machine learning, including TensorFlow Lite for mobile and embedded devices, TensorFlow.js for web applications, and

TensorFlow Serving for deployment. Widely used in both research and production environments. Uses a static computation graph.

Pytorch:

Matplotlib or other visualization libraries:

When working with Generative Adversarial Networks (GANs), visualization is crucial for understanding the training progress, inspecting generated samples, and diagnosing potential issues. Matplotlib is a popular Python library for creating static, animated, and interactive visualizations. In the context of GANs, Matplotlib can be used to visualize various aspects of the training process and the generated images. Additionally, other specialized visualization libraries can complement Matplotlib for specific tasks in GAN development. Here's an explanation:

Matplotlib in GANs:

Training Progress:

Plotting and tracking the generator and discriminator losses over time to observe convergence or potential issues. Visualizing the distribution of real and generated samples to ensure diversity and quality.

Generated Samples:

Displaying randomly generated samples at different stages of training to visually assess the quality and progression.

Creating side-by-side comparisons of real and generated images for qualitative evaluation.

Intermediate Layer Activations:

Visualizing the activations of intermediate layers in the generator and discriminator to understand feature representations.

Image Transforms:

Showing the effects of different image transformations, such as rotations or translations, on generated samples.

Other Visualization Libraries:

TensorBoard (for TensorFlow):

TensorFlow provides TensorBoard, an interactive visualization tool. It can be used to monitor training metrics, visualize the graph, and inspect the generated samples over time.

Visdom:

Visdom is a Python library that facilitates real-time interactive visualization. It's commonly used for monitoring training progress and visualizing results during GAN training.

Seaborn:

Seaborn is built on top of Matplotlib and provides a high-level interface for statistical data visualization. It can enhance the aesthetics of plots for better readability.

Plotly:

Plotly is useful for creating interactive plots. It can be employed to build 3D visualizations of generated samples or explore high-dimensional latent spaces.

OpenCV:

OpenCV is primarily an image processing library but can be useful for visualizing images, especially when dealing with pre-processing steps or augmentations in GANs.

How Visualization Supports GAN Development:

Debugging:

Identifying training issues, such as mode collapse or vanishing gradients, through visual analysis of loss curves.

Hyperparameter Tuning:

Assessing the impact of hyperparameter changes on generated samples and discriminator decisions.

Understanding Latent Space:

Visualizing the latent space and how it maps to generated images can aid in understanding the diversity of generated samples.

User Feedback:

Generating images for qualitative evaluation by users or stakeholders.

In summary, Matplotlib, along with other visualization libraries, plays a crucial role in monitoring, analyzing, and improving GANs by providing insights into the model's behavior and the quality of generated samples.

Custom scripts or notebooks:

Custom scripts and Jupyter Notebooks are essential tools when working with GAN models. They facilitate the development, training, and analysis of GANs by providing a flexible and interactive environment.

Testing frameworks (e.g., Pytest): for automating and managing test suites.

This comprehensive testing strategy ensures the robustness and reliability of the human face generation model at various stages of development, from data preparation to deployment. Regularly reviewing and updating the testing strategy enhances the model's overall performance and mitigates potential issues.

4.2 Outcomes:


Following are the outcomes for our Project :

1.



Epoch	[1/	1]	d_loss:	g_loss:
Epoch	[1/	1]	d_loss: 1.5603	g_loss: 2.0594
Epoch	[1/	1]	d_loss: 0.9731	g_loss: 2.1177
Epoch	[1/	1]	d_loss: 1.2468	g_loss: 1.9661
Epoch	[1/	1]	d_loss: 1.0558	g_loss: 1.7339
Epoch	[1/	1]	d_loss: 1.2615	g_loss: 1.1507
Epoch	[1/	1]	d_loss: 1.1214	g_loss: 1.4842
Epoch	[1/	1]	d_loss: 1.3015	g_loss: 0.8161
Epoch	[1/	1]	d_loss: 1.2894	g_loss: 1.1741
Epoch	[1/	1]	d_loss: 1.4067	g_loss: 0.9555
Epoch	[1/	1]	d_loss: 0.9958	g_loss: 1.2182
Epoch	[1/	1]	d_loss: 1.1380	g_loss: 1.0279
Epoch	[1/	1]	d_loss: 1.3749	g_loss: 2.4019
Epoch	[1/	1]	d_loss: 1.3059	g_loss: 1.7136
Epoch	[1/	1]	d_loss: 1.4158	g_loss: 2.3569
Epoch	[1/	1]	d_loss: 1.1837	g_loss: 0.9826
Epoch	[1/	1]	d_loss: 1.3492	g_loss: 2.4449
Epoch	[1/	1]	d_loss: 1.1143	g_loss: 1.1946
Epoch	[1/	1]	d_loss: 1.3758	g_loss: 0.8122
Epoch	[1/	1]	d_loss: 1.2043	g_loss: 1.6733
Epoch	[1/	1]	d_loss: 1.1825	g_loss: 1.3239
Epoch	[1/	1]	d_loss: 0.9562	g_loss: 1.4013
Epoch	[1/	1]	d_loss: 1.1063	g_loss: 1.2294
Epoch	[1/	1]	d_loss: 1.1037	g_loss: 0.8244
Epoch	[1/	1]	d_loss: 0.9094	g_loss: 0.6511
Epoch	[1/	1]	d_loss: 1.3061	g_loss: 2.5596
Epoch	[1/	1]	d_loss: 1.0546	g_loss: 0.8288
Epoch	[1/	1]	d_loss: 1.3593	g_loss: 1.3532
Epoch	[1/	1]	d_loss: 1.2538	g_loss: 0.9451
Epoch	[1/	1]	d_loss: 1.1810	g_loss: 0.9429
Epoch	[1/	1]	d_loss: 0.7384	g_loss: 1.8655
Epoch	[1/	1]	d_loss: 1.5362	g_loss: 1.8812
Epoch	[1/	1]	d_loss: 1.6230	g_loss: 1.1147
Epoch	[1/	1]	d_loss: 0.9572	g_loss: 1.2195
Epoch	[1/	1]	d_loss: 1.2418	g_loss: 0.9699
Epoch	[1/	1]	d_loss: 1.2982	g_loss: 1.4089
Epoch	[1/	1]	d_loss: 1.1597	g_loss: 1.6815
Epoch	[1/	1]	d_loss: 1.2273	g_loss: 1.8716
Epoch	[1/	1]	d_loss: 1.2789	g_loss: 1.0278
Epoch	[1/	1]	d_loss: 1.0397	g_loss: 1.2032
Epoch	[1/	1]	d_loss: 1.1574	g_loss: 0.9541
Epoch	[1/	1]	d_loss: 1.2318	g_loss: 1.5029
Epoch	[1/	1]	d_loss: 1.0733	g_loss: 1.4316
Epoch	[1/	1]	d_loss: 1.0913	g_loss: 0.9806
Epoch	[1/	1]	d_loss: 1.1013	g_loss: 1.2359
Epoch	[1/	1]	d_loss: 1.2203	g_loss: 1.4926
Epoch	[1/	1]	d_loss: 1.5319	g_loss: 1.3209

3.



4.




fig 12:- Face Generated

Chapter 5: Results and Evaluation

4.3 Results (presentation of findings, interpretation of the results, etc.)

4.4 Comparison with Existing Solutions (if applicable)

When comparing your GAN-based human face generation project with existing solutions, it's essential to consider various aspects such as model performance, training efficiency, diversity of generated faces, and ethical considerations. Here's an elaboration on how to approach the comparison:

1. Model Performance:

Metrics: Evaluate the quantitative performance of your model using metrics such as Inception Score, Fréchet Inception Distance (FID), or other relevant measures.

Comparison: Compare your model's performance with state-of-the-art GANs for face generation. Highlight any improvements or unique features in your approach.

2. Training Efficiency:

Convergence Speed: Assess how quickly your GAN converges during training compared to existing solutions.

Stability: Consider the stability of training, including resistance to mode collapse and the smoothness of the learning curve.

3. Diversity of Generated Faces:

Visual Inspection: Perform qualitative evaluation by visually inspecting the diversity of faces generated by your model.

Attribute Control: If applicable, compare the ability of your model to control specific attributes (e.g., age, gender) with existing solutions.

4. Ethical Considerations:

Bias and Fairness: Assess how well your model addresses biases related to gender, ethnicity, or other sensitive attributes. Consider fairness metrics.

Privacy: Highlight any privacy-preserving measures implemented in your model, especially if working with identifiable facial data.

5. User Interaction and Interface :

User-Friendliness: If your project involves a user interface, compare its user-friendliness with existing solutions. **Customization:** Evaluate the level of customization and control users have over the generated faces.

6. Resource Requirements:

Computational Resources: Compare the computational resources (CPU, GPU) required for training and inference with existing solutions.

Memory Usage: Assess the memory efficiency of your model during training and generation.

7. Documentation and Accessibility:

Documentation Quality: Emphasize the clarity and completeness of your project's documentation, making it easy for others to understand and use.

Open Source: If applicable, compare the openness of your project (availability of code, pre-trained models) with existing solutions.

8. Scalability:

Resolution of Generated Images: If your GAN generates high-resolution images, compare its scalability with existing solutions in terms of memory and computational requirements.

9. Innovation and Novelty:

Novel Approaches: Highlight any novel techniques, architectures, or approaches you introduced that set your project apart from existing solutions.

Contributions: If your project makes contributions to the field, such as new datasets or evaluation metrics, emphasize these contributions.

10. Limitations and Challenges:

Transparency: Be transparent about the limitations and challenges of your model. Discuss areas where improvements can be made.

Potential Biases: Acknowledge and address any potential biases or shortcomings in your model.

11. Deployment Considerations:

Ease of Deployment: Consider how easy it is to deploy your model in real-world scenarios compared to existing solutions.

Integration: If applicable, discuss the ease of integration with existing systems or frameworks.

12. Long-Term Maintenance:

Sustainability: Discuss your plans for maintaining and updating the project over time. Consider the sustainability of your solution compared to existing models.

13. Regulatory Compliance:

Compliance: Ensure that your model adheres to relevant regulatory standards, especially when dealing with sensitive data.

In summary, a comprehensive comparison with existing solutions involves evaluating your GAN model from various angles, considering both quantitative and qualitative aspects. Clearly communicate the strengths, innovations, and potential areas for improvement in your project.

Chapter 6: Conclusions and Future Scope

5.1 Conclusion (summarize key findings, limitations and contributions to the field)

The implementation of the Deep Convolutional Generative Adversarial Network (DCGAN) for human face generation has yielded noteworthy findings, while also acknowledging certain limitations and contributions to the field.

Key Findings:

Image Generation Quality:

The DCGAN has demonstrated the ability to generate realistic and high-quality human face images. The generated samples exhibit features consistent with the training data, showcasing the effectiveness of the model architecture.

Training Stability:

The stability of the training process is crucial for GANs, and the implemented DCGAN has shown resilience against common training challenges such as mode collapse. The use of convolutional layers and normalization techniques contributes to a more stable and convergent training process.

Latent Space Exploration:

The generator has successfully learned a meaningful latent space representation. This is evident in the diverse range of facial features present in the generated images, indicating that the model has captured important variations in the data.

Ethical Considerations:

Ethical considerations in face generation, such as avoiding biases and ensuring privacy, have been acknowledged. Future iterations of the model could integrate more advanced techniques to address these concerns.

Limitations:

Data Limitations:

The performance of the model is inherently tied to the quality and diversity of the training data. Limitations in the training dataset, such as insufficient diversity or size, may impact the model's

ability to generalize to a broader range of faces.

Hyperparameter Sensitivity:

GANs are sensitive to hyperparameter choices, and finding optimal settings can be a non-trivial task. Further experimentation and tuning may be required to achieve optimal results.

Ethical Considerations:

While ethical considerations have been acknowledged, ensuring fairness and mitigating biases in generated faces remain ongoing challenges. Continued research and development are needed to address these ethical concerns comprehensively.

Contributions to the Field:

Open-Source Implementation:

The provided implementation serves as an open-source resource for researchers and developers interested in GANs, particularly for human face generation. The codebase can be utilized, extended, and modified for various image synthesis tasks.

Understanding Latent Representations:

The model contributes to the understanding of latent space representations in GANs. Visualization of the latent space and its impact on generated samples provides insights into the learning process of the generator.

Training Stability Techniques:

The implemented model incorporates techniques for stabilizing GAN training, including the use of convolutional layers and batch normalization. These practices contribute to a more robust and convergent training process.

Future Directions:

Dataset Enhancement:

To address data limitations, future work could focus on acquiring a more diverse and extensive dataset for training. This could lead to improved generalization and the generation of faces with a wider range of characteristics.

Ethical Advancements:

Ongoing research should explore advanced techniques to mitigate biases and enhance the ethical

considerations of generated faces. This could involve incorporating fairness-aware training strategies and privacy-preserving mechanisms.

Hyperparameter Optimization:

Further experimentation with hyperparameter settings may uncover more optimal configurations for the model. Techniques such as automated hyperparameter tuning could be employed to streamline this process.

In conclusion, the implemented DCGAN has demonstrated promising capabilities in human face generation. While acknowledging its limitations, the model's contributions to the understanding of latent spaces and training stability make it a valuable asset in the field of generative models. Future advancements in data quality, ethical considerations, and model refinement hold the potential for even more impactful results.

5.2 Future Scope

The future scope for a Human Face Generation using GAN project is broad and holds potential for various advancements and applications. Here are several potential avenues for future development and research:

High-Resolution Image Generation:

Enhance the model to generate high-resolution facial images. This involves optimizing the architecture, training strategies, and potentially exploring progressive growing techniques to handle larger image sizes.

Improved Latent Space Manipulation:

Research and develop techniques for more intuitive and controllable manipulation of the latent space. This could involve exploring disentangled representations to control specific facial attributes independently.

Dynamic Facial Expressions:

Extend the model to generate dynamic facial expressions. This involves capturing temporal dependencies and variations to create sequences of images representing different facial expressions.

Interactive User Interfaces:

Develop interactive user interfaces that allow users to customize and interact with the generated faces. This could involve real-time manipulation of facial features or incorporating user feedback into the training process.

Cross-Domain Face Generation:

Explore the generation of faces across different domains, such as transforming sketches or artistic representations into realistic faces. This expands the applicability of the model beyond standard photographic datasets.

Biometric Applications:

Investigate the use of generated faces for biometric applications, such as face recognition or facial emotion analysis. This requires ensuring the generated faces are not only visually realistic but also functionally accurate.

Ethical and Fairness Considerations:

Invest in research that addresses ethical concerns related to biases in generated faces. Implement fairness-aware training techniques and mechanisms to ensure that the model produces diverse and unbiased results.

Data Augmentation and Privacy Preservation:

Explore techniques for privacy-preserving face generation, especially in scenarios where generating faces with specific attributes could be privacy-sensitive. Additionally, investigate methods for augmenting limited training data to improve model generalization.

Real-Time Applications:

Optimize the model for real-time applications, such as video game character generation, virtual reality environments, or video conferencing platforms. This involves considerations for inference speed and model deployment in resource-constrained environments.

References

- [1]. Jiankang Deng, Jia Guo, Niannan Xue, Stefanos Zafeiriou; Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019, pp. 4690-4699.
- [2]. Lingzhi Li, Jianmin Bao, Hao Yang, Dong Chen, Fang Wen; “FaceShifter: Towards High Fidelity And Occlusion Aware Face Swapping”; 2020, arXiv:1912.13457v3.
- [3]. Yujun Shen, Jinjin Gu, Xiaoou Tang, Bolei Zhou; Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2020, pp. 9243-9252.
- [4]. Amanda Duarte, Francisco Roldan, Miquel Tubau, Janna Escur, Santiago Pascual, Amaia Salvador, Eva Mohedano, Kevin McGuinness, Jordi Torres, Xavier
- [5]. Giro-i-Nieto; “Wav2Pix: Speech-conditioned Face Generation using Generative Adversarial Networks”; 2019, arXiv:1903.10195v1.
- [6]. Bing Li, Yuanlue Zhu, Yitong Wang, Chia-Wen Lin, Bernard Ghanem, Linlin Shen; “AniGAN: Style-Guided Generative Adversarial Networks for Unsupervised Anime Face Generation” ; 2021, arXiv:2102.12593v2.
- [7]. Mingcong Liu, Qiang Li, Zekui Qin, Guoxin Zhang, Pengfei Wan, Wen Zheng;
- [8]. BlendGAN: Implicitly GAN Blending for Arbitrary Stylized Face Generation”; 2021, arXiv:2110.11728v.
- [9]. Steven Walton, Ali Hassani, Xingqian Xu, Zhangyang Wang, Humphrey Shi; “StyleNAT: Giving Each Head a New Perspective” ; 2022, arXiv:2211.05770v2.
- [10]. Boyang Deng, Yifan Wang, Gordon Wetzstein; “LumiGAN: Unconditional Generation of Relightable 3D Human Faces” ; 2023, arXiv:2304.13153.
- [11]. Generative Adversarial Nets" by Ian Goodfellow et al. (2014)
- [12]. “Progressive Growing of GANs for Improved Quality, Stability, and Variation" by Tero

Karras et al.(2018)

- [13]. "StyleGAN: A Style-Based Generator Architecture for Generative Adversarial Networks" by Tero Karras et al. (2019)
- [14]. "StyleGAN2: Analyzing and Improving the Image Quality of StyleGAN" by Tero Karras et al. (2020)
- [15]. "Deep Residual Learning for Image Recognition" by Kaiming He et al. (2016)
- [16]. "CycleGAN: Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks" by Jun-Yan Zhu et al. (2017)
- [17]. "High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs" by Ting-Chun Wang et al. (2018)
- [18]. "GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium" by Martin Heusel et al. (2017)
- [19]. "StarGAN: Unified Generative Adversarial Networks for Multi-Domain Image-to-Image Translation" by Yunjey Choi et al. (2018)
- [20]. "BigGAN: Large Scale GAN Training for High Fidelity Natural Image Synthesis" by Andrew Brock et al. (2018)
- [21]. "Wasserstein GAN" by Martin Arjovsky et al. (2017)
- [22]. "Self-Attention Generative Adversarial Networks" by Han Zhang et al. (2019)
- [23]. "FID - Frechet Inception Distance" by Martin Heusel et al. (2018)
- [24]. "Few-Shot Adversarial Learning of Realistic Neural Talking Head Models" by Egor Zakharov et al. (2019)
- [25]. "Deep Learning" by Ian Goodfellow, Yoshua Bengio, and Aaron Courville (2016)
- [26]. "GANs in Action: Deep Learning with Generative Adversarial Networks" by Jakub Langr and Vladimír Bok (2019).

Appendix

CODE:

Commented out IPython magic to ensure Python compatibility.

```
import pickle as pkl
```

```
import
```

```
matplotlib.pyplot
```

```
plt as plt import
```

```
numpy as np
```

```
import
```

```
unittest
```

```
as tests import
```

```
helper
```

```
# %matplotlib inline
```

```
!unzip '/content/processed-
```

```
celeba-small.zip' data_dir =
```

```
'processed_celeba_small/'
```

```
import torch
```

```
from torchvision import
```

```
datasets as dset from
```

```
torchvision import
```

```
transforms
```

```
def get_dataloader(batch_size, image_size, data_dir='processed_celeba_small/):
```

```
dataset
```

```
    transform=transfo
```

```

        rms.Compose([
            transforms.Resi
            ze(image_size),
            transforms.ToT
            ensor(),
        ])

    data_loader = torch.utils.data.DataLoader(dataset, batch_size=batch_size,
                                              shuffle=True)

    return data_loader

# Define function hyperparameters

batch_size = 16

img_size = 32

# Call your function and get a dataloader

celeba_train_loader =

get_dataloader(batch_size, img_size)#

helper display function

def

    imsho

    w(img

    ):

    npimg

    =

    img.nu

    mpy()

    plt.imshow(np.transpose(n

```

```

ping, (1, 2, 0))# obtain one
batch of training images

dataiter =
iter(celeba_train_loader)

images, _ = next(iter(celeba_train_loader)) # _ for no labels

# plot the images in the batch, along with the
corresponding labelsfig = plt.figure(figsize=(20,
4))

plot_size=16

for idx in np.arange(plot_size):

    ax = fig.add_subplot(2, int(plot_size/2), idx+1,
xticks=[], yticks=[])imshow(images[idx])

def scale(x, feature_range=(-1, 1)):

    """ Scale takes in an image x and returns
that image, scaledwith a feature_range
of pixel values from -1 to 1.

This function assumes that the input x is already
scaled from 0-1."# assume x is scaled to (0, 1)

# scale to feature_range and return scaled x

return x * (feature_range[1] - feature_range[0]) + feature_range[0]

```

```

# check scaled range

# should be
close to -1
to 1img =
images[0]
scaled_img
=
scale(img)
print('Min: ',
scaled_img.min(
)) print('Max: ',
scaled_img.max(
))import torch.nn
as nn
import torch.nn.functional as F

```

CNN Layer Addition :

```

def make_conv(in_channels, out_channels, kernel_size, stride=2, padding=1,
batch_norm=True):layers=[]
conv_layer = nn.Conv2d(in_channels,
out_channels,
kernel_size, stride,
padding, bias=False)
layers.append(conv_layer)
# if batch norm set to True add a
batch norm layerif batch_norm:

```

```

        layers.append(nn.BatchNorm2d
(out_channels))return
nn.Sequential(*layers)
class
Discriminator(
nn.Module):
def __init_
(self,
conv_dim)
super(Discriminator, self)._
init_()self.conv_dim =
conv_dim
# first layer : input 32 x 32 with no batch norm
self.conv1 = make_conv(3, conv_dim, 4,
batch_norm=False)# second layer : input 16 x
16 with batch norm

```

```

self.conv2 = make_conv(conv_dim ,
conv_dim*2, 4)# third layer : input 8
x 8 with batch norm

self.conv3 = make_conv(conv_dim*2,
conv_dim*4, 4)# fourth layer : input 4
x 4 with batch norm

self.conv4 = make_conv(conv_dim*4,
conv_dim*8, 4)# fully connected layer
: one output (fake/real)

self.fc =
nn.Linear(conv_dim*8*2*
2, 1)def forward(self, x):
    out =
    F.leaky_relu(self.conv1
(x), 0.2) out =
    F.leaky_relu(self.conv2
(out), 0.2)out =
    F.leaky_relu(self.conv3
(out), 0.2)out =
    F.leaky_relu(self.conv4
(out), 0.2)# flatten
out = out.view(-1,
self.conv_dim*8*2*2)#
final output layer
o

```


u
t
=
s
el
f.
f
c
(
o
u
t)
r
et
u
r
n
o
u
t

tests.test_discriminator

(Discriminator)#This

cell is for a helper

function

def make_tconv(in_channels, out_channels, kernel_size, stride=2, padding=1,

```
batch_norm=True):layers=[]  
  
transpose_conv_layer = nn.ConvTranspose2d(in_channels, out_channels,  
                                           kernel_size, stride,  
padding, bias=False)# append transpose  
convolutional layer  
  
layers.append(transpose_conv_layer)
```

```

if batch_norm:
    layers.append(nn.BatchNorm2d
        (out_channels))

return
nn.Sequential(*l
ayers)class
Generator(nn.M
odule):
    def __init__(self, z_size,
        conv_dim):
        super(Generator, self)._
        init__() self.conv_dim =
        conv_dim
        # layers
        # first convolutional layer : input 2 x 2
        self.tconv1 = make_tconv(conv_dim*8,
        conv_dim*4, 4)#second convolutional
        layer : input 4 x 4
        self.tconv2 = make_tconv(conv_dim*4,
        conv_dim*2, 4)# third convolutional
        layer : input 8 x 8
        self.tconv3 =
        make_tconv(conv_dim*2, conv_dim,
        4)# last convolutional layer : output
        32 x 32 x 3

```

```

self.tconv4 = make_tconv(conv_dim, 3, 4,
batch_norm=False)self.fc = nn.Linear(z_size,
conv_dim*8*2*2)
#
complete
init
function
def
forward(se
lf, x):
# fully-
connected +
reshapeout
= self.fc(x)
out = out.view(-1, self.conv_dim*8, 2, 2) #
(batch_size, depth, 4, 4)# hidden transpose conv
layers + relu
out = F.relu(self.tconv1(out))
out = F.relu(self.tconv2(out))

```

```

    out =
    F.relu(self.tco
    nv3(out))#
    last layer
    out =
    self.tco
    nv4(ou
    t) #
    apply
    tanh
    activati
    onout =
    torch.ta
    nh(out)
    return
    out

tests.test_gener
ator(Generator)
def
weights_init_n
ormal(m):
    # get the class name to ensure that we initialise only for convolutional
    and linear layersclass_name = m.__class__.__name
    if hasattr(m, 'weight') and (class_name.find('Conv') != -1 or
    class_name.find('Linear') != -1):m.weight.data.normal_(0.0, 0.02)

```

```

    # set the bias term to 0 if it exists

    if hasattr(m, 'bias') and

        m.bias is not None:

            m.bias.data.zero_()

def build_network(d_conv_dim, g_conv_dim, z_size):

    Define discriminator and generator

    D = Discriminator(d_conv_dim)

    G = Generator(z_size=z_size,

conv_dim=g_conv_dim)# initialize

model weights

D.apply(weights_init_normal)

G.apply(weights_init_normal)

    return D, G

# Define

model

hyperparams

d_conv_dim

= 128

g_conv_dim = 128

z_size = 100

D, G = build_network(d_conv_dim,

g_conv_dim, z_size)import torch

# Check for a GPU

train_on_gpu =

```

```

torch.cuda.is_available()

if not train_on_gpu:

    print('No GPU found. Please use a GPU to train your
neural network.')else:

    print("Trai
ning on
GPU!")def
real_loss(D
_out):

    batch_size = D_out.size(0)

    labels = torch.ones(batch_size)*0.9 #
performed smoothingif train_on_gpu:

        labels = labels.cuda()

    criterion =
nn.BCEWithLogitsLoss

    () loss =
criterion(D_out.squeeze

    (),labels)return loss

def
fake_loss(
D_out):

    batch_size
=
D_out.size(
0)

```

```
labels = torch.zeros(batch_size) # fake images
```



```

if train_on_gpu:
    labels = labels.cuda()

criterion =
nn.BCEWithLogitsLoss

() loss =
criterion(D_out.squeeze

(),labels)return loss

import torch.optim as optim

# Create optimizers for the discriminator D
and generator Gd_lr = 0.0002

g_lr = 0.0004

d_optimizer = optim.Adam(D.parameters(),d_lr,
betas=(0.2, 0.999))g_optimizer =
optim.Adam(G.parameters(),g_lr, betas=(0.2,
0.999))def train(D, G, n_epochs, print_every=50):

    # move
    models
    to GPU

    if
    train_o
    n_gpu:

        D.cuda()

        G.cuda()

    # keep track of loss and generated,

    "fake" samplesamples = []

```

```
losses = []

# Get some fixed data for sampling. These are images that are held
# constant throughout training, and allow us to inspect the
# model's performance
sample_size=16

fixed_z = np.random.uniform(-1, 1,
                             size=(sample_size, z_size))
fixed_z =
torch.from_numpy(fixed_z).float()

# move z to GPU if available
```

```

if train_on_gpu:
    fixed_z =
fixed_z.cuda
a()# epoch
training
loop
for epoch in
    range(n_epo
        chs):# batch
        training
        loop
        for batch_i, (real_images, _) in
            enumerate(celeba_train_loader):batch_size =
            real_images.size(0)
            real_images = scale(real_images)
            # 1. Train the discriminator on real
            and fake images
            d_optimizer.zero_grad()
            if train_on_gpu:
                real_images =
                real_images.cuda()#
                loss on real images
                d_real =
                D(real_images)
                d_real_loss =

```

```
real_loss(d_rea
l)#train with
fake images
z = np.random.uniform(-1, 1,
size=(batch_size, z_size))z =
torch.from_numpy(z).float()
i
f
t
r
a
i
n
-
o
n
-
g
p
u
:
z
```

```
=  
  
z  
.  
c  
u  
d  
a  
(  
)  
fake_i  
mage  
s =  
G(z)#  
loss  
on  
fake  
image  
s  
d_fake = D(fake_images)  
d_fake_loss =  
fake_loss(d_fake  
)# backprop
```

```

d_loss =
d_real_loss +
d_fake_loss
d_loss.backward()
d_optimizer.step()
# 2. Train the generator with an
adversarial loss
g_optimizer.zero_grad()

```

Generate fake images

```

z = np.random.uniform(-1, 1,
size=(batch_size, z_size))z =
torch.from_numpy(z).float()
fake_imag
es = G(z)
d_fake =
D(fake_im
ages)
g_loss =
real_loss(d
_fake)#
perfor
backprop
g_loss.bac
kward()

```

```
g_optimize

r.step()

# Print some loss stats

if batch_i % print_every == 0:

    # append discriminator loss

    and generator loss

    losses.append((d_loss.item(),

g_loss.item())) # print

discriminator and generator

loss

print('Epoch [{:5d}/{:5d}] | d_loss: {:.6f} | g_loss:

{:.6f}'.format(epoch+1, n_epochs,

d_loss.item(), g_loss.item()))
```

```
## AFTER EACH EPOCH##
```

```
# this code assumes your generator is named G, feel free to
```

```
change the name# generate and save sample, fake images
```

```
G.eval() # for
```

```
generating
```

```
samplesamples_z
```

```
= G(fixed_z)
```

```
samples.append(s
```

```
amples_z)
```

```
G.train() # back
```

```
to training mode#
```

```
Save training
```

```
generator samples
```

```
with
```

```
open('train_samples.p
```

```
kl', 'wb') as f:
```

```
pkl.dump(samples, f)
```

```
#
```

```
finally
```

```
return
```

```
losses
```

```
return
```

```
losses
```

```
# set
```



```

number
of
epochs
n_epoc
hs = 1
# call training function
losses = train(D, G,
n_epochs=n_epochs)fig,
ax = plt.subplots()
losses = np.array(losses)
plt.plot(losses.T[0],
label='Discriminator', alpha=0.5)
plt.plot(losses.T[1], label='Generator',
alpha=0.5) plt.title("Training Losses")
plt.legend();
# helper function for viewing a list of passed in
sample imagesdef view_samples(epoch,
samples):

```

```

fig, axes = plt.subplots(figsize=(16,4), nrows=2, ncols=8,
sharey=True, sharex=True)for ax, img in zip(axes.flatten(),
samples[epoch]):
    img =
    img.detach().cpu().
    numpy()img =
    np.transpose(img,
    (1, 2, 0))
    img = ((img + 1)*255 /
    (2)).astype(np.uint8)
    ax.xaxis.set_visible(False)
    ax.yaxis.set_visible(False)
    im = ax.imshow(img.reshape((32,32,3)))
# Load samples from generator,
taken while trainingwith
open('train_samples.pkl', 'rb') as f:
    samples = pkl.load(f)
_ = view_samples(-1, samples)

```

JAYPEE UNIVERSITY OF INFORMATION TECHNOLOGY, WAKNAGHAT

PLAGIARISM VERIFICATION REPORT

Date:

Type of Document (Tick): PhD Thesis M.Tech Dissertation/ Report B.Tech Project Report Paper

Name: _____ Department: _____ Enrolment No _____ Contact No. _____

letters):

UNDERTAKING

I undertake that I am aware of the plagiarism related norms/ regulations, if I found guilty of any plagiarism and copyright violations in the above thesis/report even after award of degree, the University reserves the rights to withdraw/ revoke my degree/report. Kindly allow me to avail Plagiarism verification report for the document mentioned above.

Complete Thesis/Report Pages Detail:

- Total No. of Pages =
- Total No. of Preliminary pages =
- Total No. of pages accommodate bibliography/references =

(Signature of Student)

FOR DEPARTMENT USE

We have checked the thesis/report as per norms and found **Similarity Index** at (%). Therefore, we are forwarding the complete thesis/report for final plagiarism check. The plagiarism verification report may be handed over to the candidate.

(Signature of Guide/Supervisor)

Signature of HOD

FOR LRC USE

The above document was scanned for plagiarism check. The outcome of the same is reported below:

Copy Received on	Excluded	Similarity Index (%)	Generated Plagiarism Report Details (Title, Abstract & Chapters)	
	<ul style="list-style-type: none">• All Preliminary Pages• Bibliography/Images/Quotes• 14 Words String		Word Counts	
Report Generated on			Character Counts	
		Submission ID	Total Pages Scanned	
			File Size	

Checked by
Name & Signature

Librarian

Please send your complete thesis/report in (PDF) with Title Page, Abstract and Chapters in (Word File) through the supervisor at plagcheck.juit@gmail.com