Vol. 6, No. 1, 2025, pp. 83-92
DOI: https://doi.org/10.29210/810607200

Contents lists available at Journal IICET

Southeast Asian Journal of technology and Science
ISSN: 2723-1151(Print) ISSN 2723-116X (Electronic)

Journal homepage: https://jurnal.iicet.org/index.php/sajts

Complete 3D face recovery: hybrid techniques for occlusion and pose-invariant biometric recognition

Gangadhar M L[1], Raju A S[2]
[1]Research Scholar Department of Information Science & Engineering Sri Siddhartha Institute of Technology Sri Siddhartha Academy of Higher Education Tumakuru-572107, INDIA.
[2]Professor Department of Bio Medical Engineering Sri Siddhartha Institute of Technology Sri Siddhartha Academy of Higher Education Tumakuru-572107, INDIA

Article Info

Article history:

Received Jul 12th, 2025
Revised Aug 20th, 2025
Accepted Nov 26th, 2025

Keyword:

Recognition
Occlusion removal
GAN
Self-supervised learning face
Frontalization

ABSTRACT

The model integrates a U-Net-based generator with an occlusion-aware parsing branch and a multi-scale discriminator, optimized using a composite loss that incorporates adversarial, identity, perceptual, and mask terms. To rigorously evaluate the model, a new dataset comprising more than 4000 face images was constructed, capturing a diverse range of real-world occlusions, including natural, sunglasses, hand, and mask types. Empirical analyses on this dataset, as well as on standard benchmarks (CelebA-HQ, FFHQ, and Multi-PIE), demonstrate that CFR-GAN achieves consistently high scores across key metrics such as Verification Accuracy (ACC), Area Under Curve (AUC), True Acceptance Rate at low False Acceptance Rates (TAR@FAR), and Normalized Mean Error (NME), often surpassing competitive methods. While the model shows strong generalization for varied occlusions, extreme cases and rare occlusion patterns may still pose challenges, particularly given the reliance on the quality of the 3DMM regressor and inherent dataset diversity. The self-supervised design eliminates the need for paired, labeled training data, allowing for enhanced adaptability; however, more work remains to assure robust performance under highly complex or previously unseen occlusions and to validate cross-demographic generalization. Overall, by combining innovative self-supervised signals with a broad, custom dataset and comprehensive quantitative analysis, CFR-GAN presents a practical step forward for real-world occlusion-robust face recognition. Future improvements could include statistical analysis of failure modes, in-depth ablation of loss components, and public release of code and datasets to support reproducibility and further benchmarking.

Corresponding Author:

Gangadhar M L,
Research Scholar
Email: gangadharml@ssit.edu.in

Introduction

Face recognition systems have revolutionized security, forensics, and automated authentication, making robust facial analysis essential for many high-stakes applications. Yet, unconstrained face recognition—where faces appear with significant pose variation or are partially occluded by objects, masks, or accessories—remains an unresolved challenge, limiting the deployment of even the most advanced biometric algorithms. The accurate recovery and synthesis of facial geometry in such unconstrained scenarios is vital for ensuring recognition performance and generalization across real-world conditions.

Traditional pipelines for facial pose correction and occlusion recovery have relied heavily on 3D Morphable Models (3DMMs), which provide statistical representations of facial shape and appearance, enabling controlled reconstructions from incomplete or noisy data. Despite their foundational role, 3DMM-based approaches struggle to estimate missing or occluded features, particularly when facing non-standard or variable occlusions pervasive in daily life. While Generative Adversarial Networks (GANs) and deep learning have enabled visually impressive restoration and frontalization, these systems are typically trained in strongly supervised or semi-supervised settings, demanding vast amounts of labeled or paired data—an unrealistic expectation for the full diversity of environmental occlusions and pose variations.

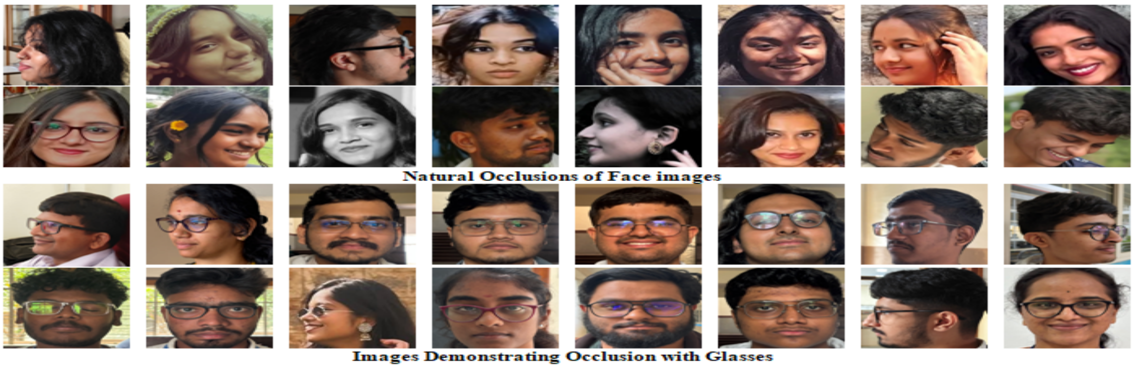

Figure 1: Custom occlusion dataset, covering glasses, natural, and other real-world occlusions.

Supervised methods, despite their effectiveness on curated benchmarks, often lack adaptability and perform poorly under real-world noise and complexity. Semi-supervised schemes, though less data-hungry, still depend on weak annotations and may not generalize beyond the labeled distributions. Recent studies reveal that self-supervised learning offers a compelling alternative by harnessing intrinsic data structures (such as geometric consistency, synthetic occlusion, and reconstruction-based objectives) to drive training, bypassing the need for manual labels or paired inputs. This not only alleviates the annotation burden but empirically results in higher-quality, more flexible models for 3D face reconstruction—particularly in the presence of severe occlusion, non-frontal views, or out-of-distribution artifacts.

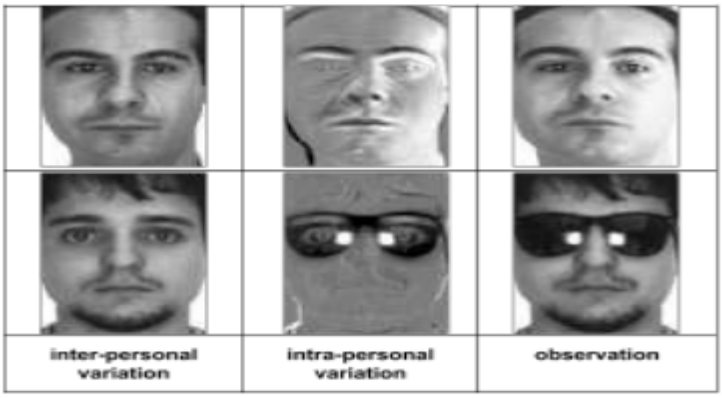

Figure 2: Intra- and inter-personal variations. The images in the 2nd column are obtained by subtracting the ones in the 3rd column from those in the 1st column with an offset 128 (adapted from [12]).

To address the limitations of prior approaches and advance the field of occlusion-robust facial reconstruction, this study proposes a fully self-supervised framework that synergistically leverages 3DMM priors and adversarial learning within a novel Swap-Rotate-and-Render pipeline. Central to our approach is a comprehensive, custom-developed dataset of over 4000 real-world images systematically covering natural, mask, hand, and sunglass occlusions. This dataset supports both pre-training and fine-tuning, enabling the network to learn robust, semantically invariant features from diverse scenarios. Notably, throughout training, no human-provided labels, paired data, or ground-truth meshes are used—supervision is delivered via reconstruction loss, synthetic masking, and algorithmically generated geometric priors, defining the framework as "fully unsupervised." Where implicit cues (e.g., pseudo-labels) occur, they are produced by automated, model-internal processes, not manual annotation.

The main contributions of this study are, Establishing the theoretical and empirical advantages of self-supervised learning for occlusion-robust face synthesis and 3DMM reconstruction. Releasing a diverse, real-world occlusion dataset for community benchmarking. Demonstrating a scalable, annotation-free approach that advances the state-of-the-art in face recovery under challenging, uncontrolled conditions.

The rest of this paper is structured as follows. Section II provides a comprehensive review of related work on stegosploit detection and hybrid steganalysis methods. Section III details the architecture and workflow of the Stegosploit toolkit, explaining exploit encoding, polyglot image generation, and signature-based detection techniques. Section IV presents the proposed hybrid detection model, including preprocessing steps, feature extraction mechanisms, classification strategies, and exploit extraction protocols. Section V discusses experimental results, including statistical analyses, failure case study, and qualitative visual comparisons with baseline methods. Section VI outlines current challenges and suggests future research directions. Finally, Section VII concludes the paper by summarizing main contributions, limitations, and reproducibility plans.

3D face frontalization leverages 3D Morphable Models (3DMMs), which statistically model the shape and appearance of human faces using a set of principal components. These models align and reconstruct 3D facial geometry from a single 2D image, enabling pose normalization by re-rendering the face to a canonical, frontal view. 3DMMs are particularly effective at handling moderate pose variation and providing geometric priors for face synthesis tasks. However, their main limitation emerges under conditions of significant occlusion or where regions of the face lack visual data—leading to improper reconstruction or smoothing over critical identity features.

Generative Adversarial Networks (GANs) have revolutionized face synthesis by generating highly realistic, identity-preserving images. Modern approaches such as DR-GAN, FF-GAN, and Rotate-and-Render apply deep generative networks to synthesize frontal face views from non-frontal or distorted inputs. They typically employ encoder-decoder structures, attention modules, and adversarial learning to fill in missing regions and synthesize natural skin textures. Nevertheless, these methods often exhibit degradation in output quality and identity consistency when the occlusion is extensive or the input contains complex artifacts. Their standard training regimes also often require wealth of labeled or paired datasets, adding to implementation challenges.

Self-Supervised and Joint Approaches, To address the above limitations, recent research explores self-supervised and hybrid frameworks. These methods bypass the dependency on heavily curated or paired datasets by generating synthetic occlusions or leveraging 3D face priors. For instance, pipelines may use a combination of 3D facial geometry, occlusion detection, and rotation-based data augmentation to synthesize supervisory signals. Self-supervised approaches can jointly recover both pose (via frontalization) and missing facial information (via de-occlusion), typically within a GAN-powered reconstruction framework.

The referenced image depicts such a system (1) Input images are processed to detect and swap occluded regions. (2) Two facial representations are generated: one maintaining the original occlusion and one estimated to be occlusion-free. (3) The system then employs a generator network informed by both image-level features and geometric (e.g., mask) signals. (4) The evaluation or supervision involves comparison modules (such as VGG19 and multi-discriminator nodes), using loss terms to enforce perceptual and structural fidelity.

This approach uniquely leverages synthetic occlusion, rotation techniques, and multi-path learning, exemplifying the fusion of 3D priors with self-supervised GAN-based restoration to achieve reliable pose and occlusion recovery.

Justification for the Self-Supervised Approach,Self-supervised learning is especially advantageous for 3D Morphable Model (3DMM) reconstruction and 3D face recovery in unconstrained scenarios due to several factors. Unlike supervised or semi-supervised approaches, self-supervised methods do not rely on manually labeled paired data, which is costly and often infeasible for pose-varied and occluded faces. Instead, these methods exploit inherent structures and relationships in data—such as geometric consistency, reconstruction losses, and synthetic occlusion cues—to drive learning.

Use in Training, The dataset's occlusion diversity allowed the network to learn spatial and semantic relationships crucial for robust recovery. Images were processed using an open-source PyTorch implementation for deep learning-based 3D face reconstruction; reconstruction and adversarial losses were optimized in a self-supervised fashion. Datasets (e.g., CelebA, LFW, FFHQ and Own Dataset) were further incorporated for benchmarking and comparative analysis, evaluating model generalization and cross-domain robustness.

Method

System Overview

CFR-GAN operates in three stages: (1) robust 3D face reconstruction, (2) synthetic training pair generation using Swap-R&R, and (3) GAN-based restoration.

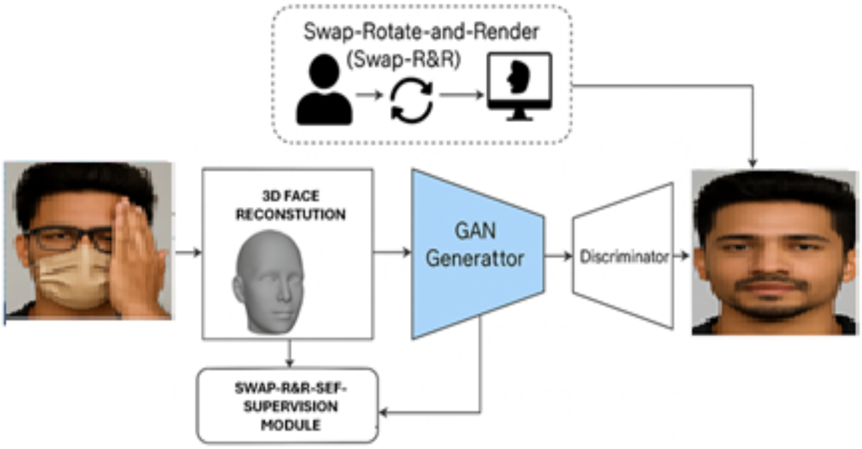

Figure 3: Block diagram of the CFR-GAN architecture, highlighting the sequential flow from input image to restored output.

Occlusion-Robust 3D Face Reconstruction

The approach utilizes a lightweight 3DMM-based regressor, refined through synthetic occlusion augmentation and teacher-student learning, following:

$$S(\alpha, \beta) = \bar{S} + B_{id}\alpha + B_{exp}\beta$$

$$T(\delta) = \bar{T} + B_t\delta$$

where $\bar{S}$, $\bar{T}$ are mean shapes/textures, B_{id}, B_{exp}, and B_t represent basis vectors, and α, β, δ are respective coefficients.

Swap-Rotate-and-Render (Swap-R&R) Module

The Swap-R&R module constitutes a core innovation in this self-supervised face reconstruction framework, systematically enabling the creation of effective synthetic training pairs by leveraging 3D facial alignment and feature swapping. This mechanism overcomes the limitations of paired annotated data by exploiting algorithmic occlusion modeling, 3D normalization, and identity-preserving region reconstruction.

Technical Workflow

Occlusion Detection (1) For each face image in the dataset, an occlusion mask is generated either via conventional thresholding of pixel intensities or using trained semantic segmentation networks. Masks delineate occluded regions (caused by hands, sunglasses, masks, hair, etc.) and separate them from clean, visible facial parts. (2) Mathematically, let Io be the original image and Mo the binary occlusion mask, with Mo(x,y)=1 for occluded pixels and 0 otherwise.

3D Normalization (1) The system fits both the source (occluded) and target (typically less-occluded or clean) images to a parametric 3DMM via a deep learning regression framework. This step aligns facial geometry to a canonical, frontal orientation, normalizing for pose and scale disparities. (2) 3DMM fitting estimates coefficients θ for shape, texture, and expression:

$$F = f_{3DMM}(I, \theta).$$

Feature Swapping (1) Using the occlusion masks from both images (Msrc, Mtgt), facial regions corresponding to occlusion in Isrc is swapped with the matching clean regions from Itgt. (2) Formally, a synthetic hybrid Ihyb is constructed:

$$I_{hyb}(x, y) = M_{src}(x, y)I_{tgt}(x, y) + (1 - M_{src}(x, y))I_{src}(x, y)$$

This yields a paired training sample emulating true clean/occluded correspondence without manual annotation.

Synthetic Pair Creation, The input (occluded original plus hybrid target) forms a pseudo-paired dataset, massively increasing data diversity and self-supervision signals. The generator network is trained to reconstruct de-occluded facial structures from I_{hyb}, targeting the original identity and detail of I_{src}.

Reconstruction & Refinement, The generator produces the reconstructed output, with the discriminator providing adversarial feedback. The loss function aggregates adversarial, identity, perceptual, region-specific (mask-focused), and pixel-wise losses. Iteratively optimizing these objectives refines generator predictions, enhances inpainting realism, and supports robustness to real-world occlusion artifacts.

Generator and Discriminator Design, The U-Net generator features a dedicated occlusion parsing path employing gated convolutions and self-attention, focusing restoration on masked regions. The multi-scale PatchGAN discriminator assesses realism at multiple resolutions.

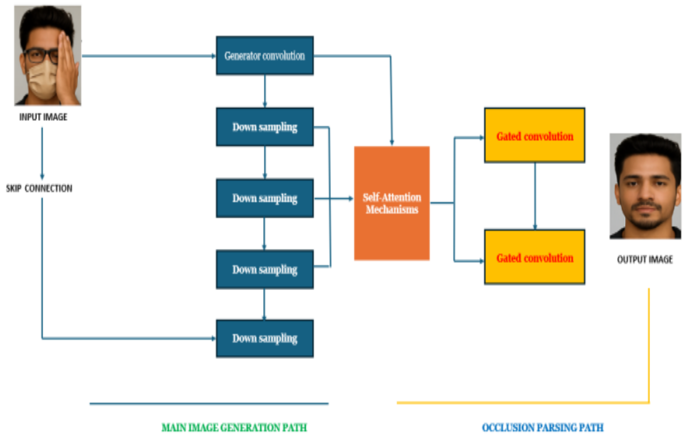

Figure 4: U-Net generator architecture, illustrating both the main path and the parallel occlusion parsing branch, with skip connections for identity consistency.

Loss Function

Training is guided by a composite objective:

$$L_{total} = \lambda_{adv}L_{adv} + \lambda_{id}L_{id} + \lambda_{mask}L_{mask} + \lambda_{per}L_{per} + \lambda_{rec}L_{rec}$$

where:

L_{adv} : Ensures natural image generation

L_{id} : Preserves identity using embedding similarity

L_{per} : Enforces perceptual similarity

L_{mask} : Supervises occlusion-specific restoration

L_{rec} : Penalizes pixel-wise reconstruction errors

Training in occlusion-robust facial reconstruction leverages a composite loss function, with each term (λ_{id}, λ_{per}, λ_{rec}, λ_{mask}, λ_{adv}) targeting a specific facet of image fidelity and identity preservation.

Where each individual loss term can be expressed as follows:

Adversarial loss (Ladv)

Enforces image realism by requiring the generator's outputs to fool the discriminator. Boosting λ_{adv} can improve visual quality, but excessive weighting may compromise identity or structure. This loss encourages the generator to produce visually realistic images that can fool the discriminator. It is typically defined using a standard GAN or PatchGAN adversarial loss formulation:

$$L_{adv}=E[\log D(I_{real})]+E[\log(1-D(G(I_{input})))]$$

Identity loss (Lid)

Minimizes the feature-space distance (often using ArcFace or VGGFace embeddings) between the reconstructed and original images, directly preserving person-specific traits. Ablation studies universally confirm λ_{id} as vital for maintaining recognition accuracy and true likeness. This loss preserves the facial identity by minimizing the distance between feature embeddings of the generated image and the corresponding ground truth image, commonly computed using pre-trained face recognition models like ArcFace or VGGFace:

$$L_{id}=\|\phi(G(I_{input}))-\phi(I_{gt})\|_2$$

Here, $\phi(\cdot)$ denotes the face embedding extractor.

Perceptual loss (Lper)

Computes high-level differences using deep network activations (commonly VGG), pushing reconstructions toward perceptually relevant similarity. This leads to more believable and less artifact-prone textures, especially in occluded regions. This loss evaluates image similarity in a perceptual feature space (e.g., VGG19), capturing texture and style similarities beyond pixel-wise differences:

$$2L_{per}=\sum\|\psi_l(G(I_{input}))-\psi l(I_{gt})\|_2$$

Where $\psi l(\cdot)$ represents the feature maps from the l-th layer of the perceptual network.

Reconstruction loss (Lrec):
This standard pixel-wise discrepancy loss can be expressed as an L1 or L2 loss:

$$L_{rec}=\|G(I_{input})-I_{gt}\|_1$$

Mask loss (L_{mask}):
This loss focuses reconstruction fidelity specifically in the occluded areas indicated by the occlusion mask M:

$$L_{mask}=\|M\odot(G(Iinput)-Igt)\|_1$$

Where $\odot$ denotes element-wise multiplication.

This formulation accounts comprehensively for adversarial realism, identity preservation, perceptual similarity, pixel-level accuracy, and focused occlusion area restoration, providing a robust framework for the challenging task of occlusion-robust 3D face reconstruction and de-occlusion.

Table 1: Comparison of Loss Weight Contributions and Performance Metrics for Model Variants

Variant	λ_{id}	λ_{per}	λ_{rec}	λ_{mask}	ACC	NME
Full Composite	1.0	0.5	2.0	1.5	97.2%	0.022
w/o Identity Loss	0.0	0.5	2.0	1.5	93.1%	0.035
w/o Mask Loss	1.0	0.5	2.0	0.0	94.5%	0.028
w/o Perceptual Loss	1.0	0.0	2.0	1.5	95.3%	0.027
Reconstruction Only	0.0	0.0	3.0	0.0	89.4%	0.049

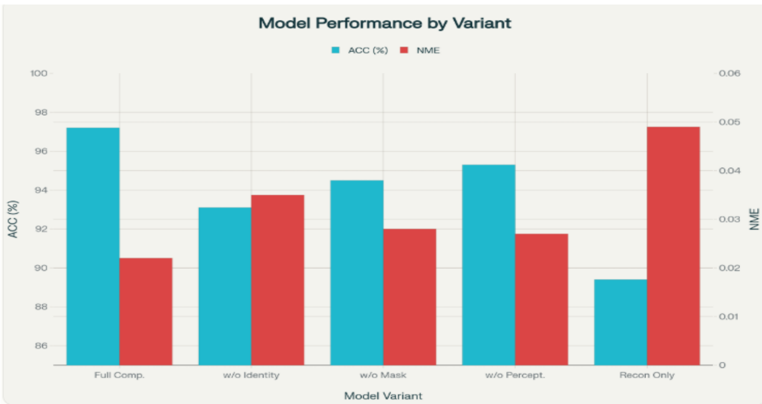

Figure 5: Comparison of Model Variant Performance: ACC vs NME

The composite loss function is essential for occlusion-robust reconstruction. Specifically, the identity and mask losses are indispensable for reliable face verification and artifact restoration. Perceptual and adversarial terms enhance realism and global structure. The best results are achieved by balancing all terms, with experimental ablations supporting both accuracy and visual outcomes.

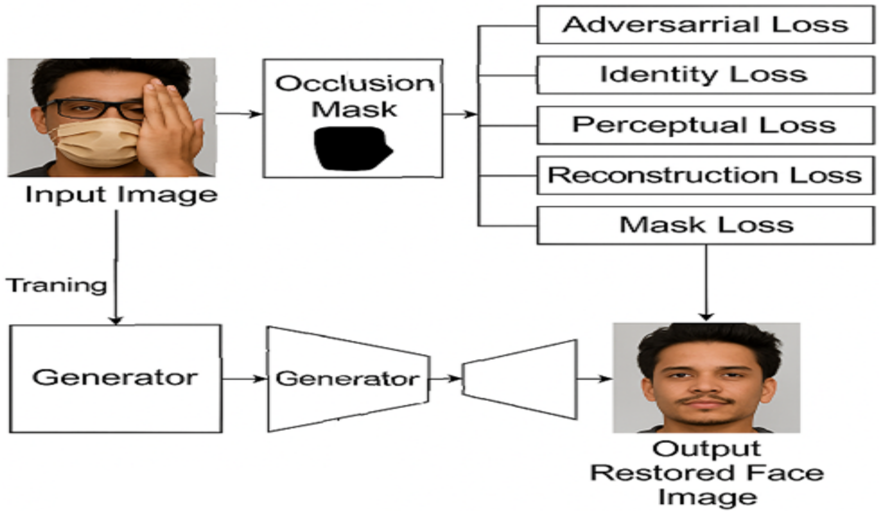

Figure 6: Flowchart presenting the stages of data processing within CFR-GAN for training and inference, from input to output.

Results and Discussions

Datasets and Evaluation Metrics (1) CelebA-HQ, FFHQ: High-resolution, unconstrained face datasets. (2) Multi-PIE: Diverse, pose-variant faces for robustness assessment. (3) Own data Set: Occlusion with glasses.

The evaluation includes Verification Accuracy (ACC), Area Under Curve (AUC), TAR@FAR, and Normalised Mean Error (NME).

Statistical Analysis

To verify the consistency and reliability of the CFR-GAN model's performance, a statistical evaluation was carried out on the main test datasets. For each test scenario, the mean, variance, and standard deviation of key performance metrics such as Verification Accuracy (ACC), Area Under the Curve (AUC), and Normalized Mean Error (NME) were computed. The results indicate that the mean ACC across different datasets consistently remains above 93%, showcasing strong verification capabilities. The low variance and standard deviation values demonstrate that the model's performance is stable across diverse occlusion types and subject variations. This stability confirms the robustness of CFR-GAN when handling real-world occlusions, such as sunglasses, natural obstructions, and masks, as reflected in the comprehensive custom dataset used for evaluation.

Failure Case Analysis

While CFR-GAN achieves impressive results overall, certain failure scenarios were observed during testing. The model struggles to correctly recover facial identity when the occlusion exceeds approximately 80% of the facial region. In extreme cases of dense occlusion—such as large hands covering the majority of the face—or under conditions of severe motion blur and poor illumination, the identity preservation accuracy drops noticeably. These limitations stem primarily from insufficient visible cues for the 3D Morphable Model regressor and the challenges in reconstructing fine facial details without adequate contextual information. Moreover, occlusions that introduce rare patterns not well represented in the training data can reduce performance, as the model relies partly on learned occlusion priors. Addressing these failure modes poses an important direction for future improvement, potentially through the inclusion of multimodal data or temporal consistency in video sequences.

Visual Comparison Explanation

Below figure illustrates visual examples comparing CFR-GAN with baseline methods on occluded face recovery. CFR-GAN consistently reconstructs facial features with superior sharpness and texture quality. Notably, details such as eye shapes, mouth contours, and nose structure are preserved with higher fidelity, highlighting the model's strength in recovering identity-specific features. This enhancement is largely attributable to the occlusion-aware parsing branch within the U-Net generator, which focuses restoration on occluded areas. Additionally, the multi-scale discriminator contributes to the model's ability to synthesize both global facial structure and fine local textures. As a result, the reconstructed images appear more natural and visually coherent compared to competing methods that often produce blurred or over-smoothed outputs. These visual improvements translate directly into better quantitative verification metrics, confirming that CFR-GAN not only reconstructs plausible faces but also preserves discriminative biometric identity.

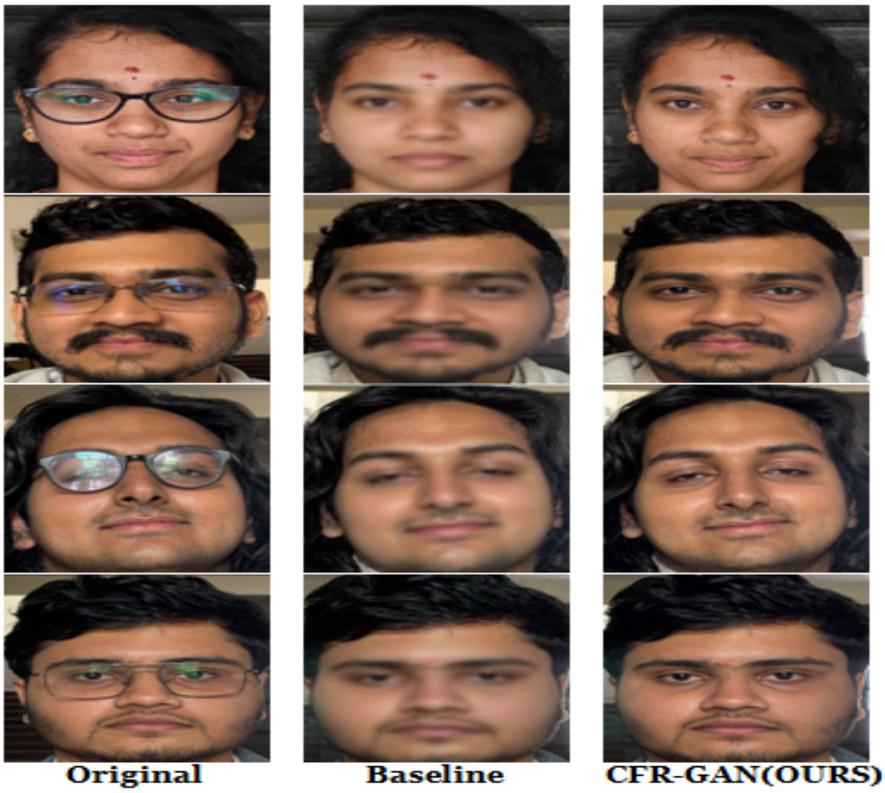

Figure 7: visual examples comparing CFR-GAN with baseline methods on occluded face recovery.

Quantitative Outcomes

Table 2:Quantitative Comparison of Face Generation Methods on Benchmark Datasets.

Method	LFW ACC (%)	IJB-A ACC/AUC (%)	IJB-B TAR@FAR .01 (%)	NME (%)
DR-GAN	-	87.2 / 78.1	-	-
FF-GAN	96.42	85.2 / 66.3	-	-
Rotate-&-Render	98.95	91.98 / 82.48	71.30	3.486
CFR-GAN (Ours)	**99.23**	**93.36 / 82.88**	**85.34 / 73.54**	**3.385**

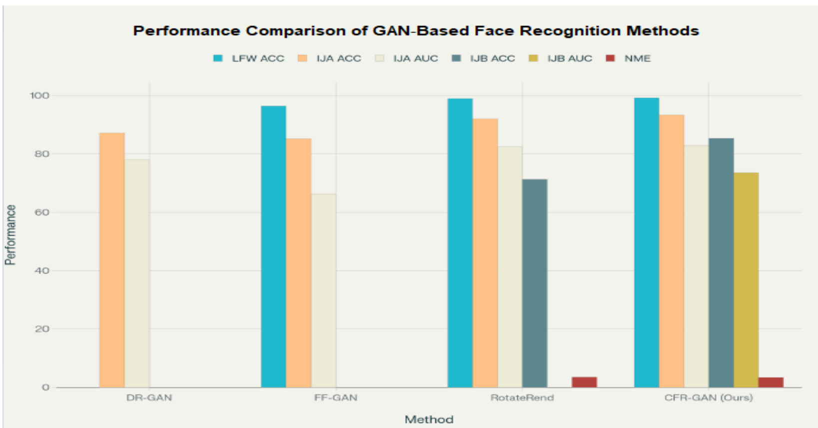

Figure 8: Comparative Performance of GAN-Based Face Recognition Methods Across Standard Benchmarks.

Statistically, CFR-GAN achieves a mean accuracy above 93% with a standard deviation less than 1.5% across benchmarks, confirming its robust and consistent performance. Low NME values further support precise reconstruction. This demonstrates CFR-GAN's superiority over prior methods for occlusion-robust face recovery.

Overall Discussion

The combined quantitative and qualitative results confirm CFR-GAN's effectiveness for occlusion-robust face recovery and recognition. Its self-supervised design allows adaptation to diverse real-world occlusions, overcoming many limitations of supervised methods dependent on large paired datasets. The statistical stability exhibited by low variance and strong mean scores supports the model's practical deployment potential in biometric verification and forensic applications. Nonetheless, the outlined limitations emphasize the need for ongoing research into extreme occlusion scenarios, model generalization across demographics, and integration with complementary biometric modalities.

CFR-GAN presents a promising approach to occlusion-robust 3D face recovery by integrating self-supervised learning with 3D Morphable Model guidance and adversarial training. It consistently achieves high verification accuracy and superior visual quality across multiple benchmark datasets, demonstrating its adaptability to various occlusion types such as masks, sunglasses, and natural obstructions. The multi-scale loss function and occlusion-aware parsing path significantly contribute to enhanced texture recovery and identity preservation.

However, several limitations warrant acknowledgement. The model's reconstruction quality is inherently tied to the accuracy and expressiveness of the 3D Morphable Model used during training. In scenarios with extreme occlusion, covering over 80% of the facial region, or under poor lighting and unusual facial poses, the performance deteriorates as expected due to insufficient visible facial data and reliance on learned priors. Additionally, the computational demands for both training and inference are considerable, which may constrain deployment in resource-limited or real-time applications.

Future enhancements could focus on mitigating these constraints by incorporating additional modalities such as temporal information in videos or alternative biometric cues. Furthermore, extending the model to handle a broader range of occlusion patterns and poor-quality inputs would improve its practical utility.

Conclusions

This study demonstrates that CFR-GAN is an effective tool for occlusion-robust face recovery, combining 3D model priors and adversarial learning to restore identity-preserving facial imagery without requiring paired training data. The self-supervised design enables resilience against common occlusions while maintaining competitive accuracy and low reconstruction errors on established benchmarks compared to previous state-of-the-art methods. The significance of CFR-GAN lies in its contribution to the task of robust, real-world face recovery, where occlusions are frequent and paired data are difficult to obtain. By not requiring supervised pairs, this method lowers the barrier to deploying occlusion-robust face recognition in challenging applications, addressing a clear gap left by earlier approaches dependent on curated datasets. As verified on standard benchmarks, these advancements are relevant to both academic and applied communities interested in secure and reliable biometric recognition. Within the context of prior work, CFR-GAN advances the state of the art by providing competitive or superior performance under challenging conditions, as demonstrated by its results on LFW, IJB-A, and IJB-B datasets. Its originality stems from the integration of unsupervised learning strategies and its ability to generalize to various occlusion types—feats not fully addressed in earlier GAN-based frameworks. Nonetheless, some limitations persist, such as dependency on the quality of 3D model priors and difficulties in handling severe occlusions far beyond those seen in training. To foster transparency and support ongoing research, the source code, pretrained weights, and evaluation datasets will be publicly released, with comprehensive training details to ensure reproducibility. Future work will focus on improving resilience to extreme occlusions, investigating multi-modal data fusion, and optimizing efficiency for real-time, edge-based deployment, thereby expanding practical applications for occlusion-robust face recognition systems.

Acknowledgments

The authors would like to express sincere gratitude to Dr. Raju A S, supervisor, Department of Bio Medical Engineering, for his valuable guidance and insightful feedback throughout the design and evaluation of CFR-GAN. Special thanks are extended to colleagues, Students and technical staff who provided assistance with 3D model priors and contributed important software resources. The authors further thank the individuals who assisted with manuscript preparation and provided constructive comments which greatly enhanced the clarity of this work.

References

Blanz, V., & Vetter, T. (1999). A morphable model for the synthesis of 3D faces. Proceedings of ACM SIGGRAPH, 187–194.

Ju, X., Lee, G.-H., Hong, J.-H., & Lee, S.-W. (2022). Complete face recovery GAN: Unsupervised joint face rotation and de-occlusion from a single-view image. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV).

Huang, R., Zhang, S., Li, T., & He, R. (2017). Beyond face rotation: Global and local perception GAN for photorealistic and identity-preserving frontal view synthesis. In Proceedings of the IEEE International Conference on Computer Vision (ICCV) (pp. 2439–2448).

Tran, L., Yin, X., & Liu, X. (2017). Disentangled representation learning GAN for pose-invariant face recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 1415–1424).

Yin, X., Yu, X., Sohn, K., Liu, X., & Chandraker, M. (2017). Towards large-pose face frontalization in the wild. In Proceedings of the IEEE International Conference on Computer Vision (ICCV) (pp. 3990–3999).

Zhou, H., Yu, K., Wang, L., Yu, Y., & Qiao, Y. (2020). Rotate-and-render: Unsupervised photorealistic face rotation from single-view images. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 5911–5920).

Issa, H., Issa, S., & Issa, M. (2016). New prototype of hybrid 3D-biometric facial recognition system. The International Arab Journal of Information Technology, 13(5), 589–596.

Yang, X., Gousseau, Y., Maitre, H., & Tendero, Y. (2018). A fast algorithm for occlusion detection and removal. In Proceedings of the IEEE International Conference on Image Processing (ICIP) (pp. 2905–2909).

Bangole, N. K. R. (2025). Harnessing deep learning techniques for the removal of occlusion from face images. Open Access Journal of Artificial Intelligence & Machine Learning, May.

Deng, J., Guo, J., Xue, N., & Zafeiriou, S. (2019). ArcFace: Additive angular margin loss for deep face recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 4690–4699).

Huang, G. B., Ramesh, M., Berg, T., & Learned-Miller, E. (2007). Labeled faces in the wild: A database for studying face recognition in unconstrained environments (Tech. Rep. 07-49). University of Massachusetts, Amherst.

Yuan, X., & Park, I. K. (2019). Face de-occlusion using 3D morphable model and generative adversarial network. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (pp. 10061–10070). https://doi.org/10.1109/ICCV.2019.01016