

Highlights

Problem Statement:

Given a face image with arbitrary pose, generate a frontal face of the same identity.

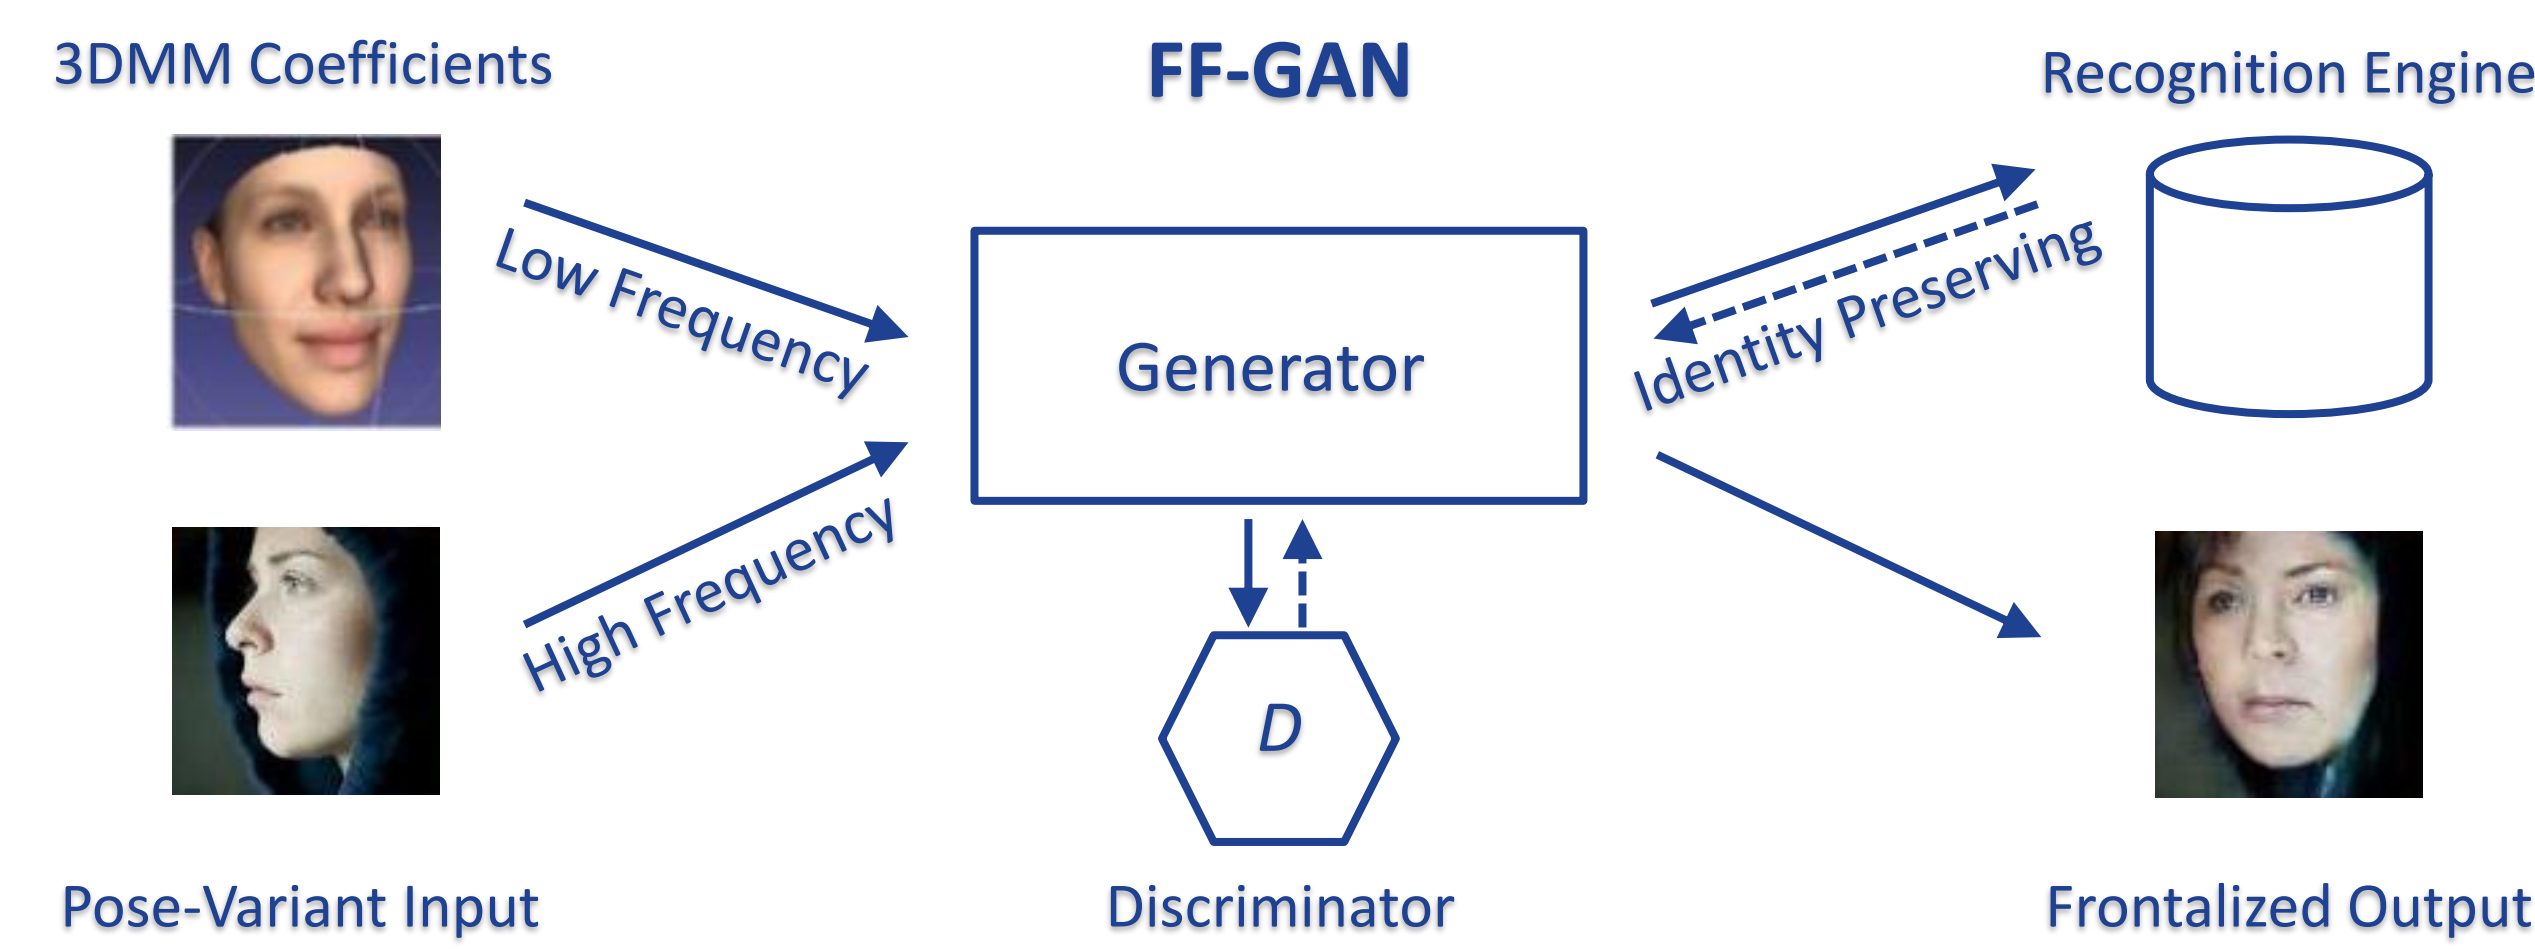
Need:

Face recognition engines favor frontal poses due to dataset biases.

Augmented reality requires photorealism from arbitrary viewpoints.

Insights and Contributions:

- ✧ Faces are constrained shapes: 3DMM priors.
- ✧ Low frequency bias in reconstruction: Adversarial framework.
- ✧ Special properties of faces: Smoothness and symmetry.
- ✧ Identity preservation: Face recognition engine.



Preliminaries

(1) 3D Morphable Model (3DMM)

3DMM defines face shape and texture in the PCA space.

Weak perspective projection: pitch, yaw, roll, scale, x-y-translations.

3DMM coefficients: shape and texture basis + projection matrix.

(2) Generative Adversarial Network (GAN)

GAN maps from a source distribution to a target distribution using a minimax optimization between a generator and a discriminator.




(3) Face Frontalization


Prior work [1]: use a single 3D surface as an approximation of any face shape.

Prior work [2]: 3DMM-based pose and expression normalization.

Drawbacks: small pose, artifacts.



FF-GAN



Training Set:

300W-LP: 122,450 images augmented from 300W.

For each training sample:

- Face image pair: x and x^g
- 3DMM coefficients: p^g



Module	Input	Output	Loss
R	x	p	weighted L2 between p and p^g
G	x and p	x^f	L1 reconstruction, smoothness, symmetric; fool D to classify x^f as real; fool C to classify x^f as the same identity
D	x^g or x^f	real / generated	cross-entropy classification loss
C	x or x^f	identity	cross-entropy classification loss

Ablation Study




Quantitative results of ablation study on Multi-PIE.

Removed module	-	C	D	R	G_{id}	G_{tv}	G_{sym}
Performance (syn.)	74.2	59.2	73.4	68.5	69.3	72.9	73.1

Experimental Results

(1) 3D Face Reconstruction



(1) Face Recognition

Face verification results on LFW, compared to face frontalization methods.

	ACC (%)	AUC (%)
Ferrari et al.	-	94.29
Hassner et al. [1]	93.62 ± 1.17	98.36 ± 0.06
HPEN [2]	96.25 ± 0.76	99.39 ± 0.02
FF-GAN (syn.)	96.42 ± 0.89	99.45 ± 0.03

Face recognition results on IJB-A.

	Verification		Identification	
	FAR=.01	FAR=.001	Rank-1	Rank-5
Wang et al.	72.9 ± 3.5	51.0 ± 6.1	82.2 ± 2.3	93.1 ± 1.4
DCNN	78.7 ± 4.3	-	85.2 ± 1.8	93.7 ± 1.0
DR-GAN	77.4 ± 2.7	53.9 ± 4.3	85.5 ± 1.5	94.7 ± 1.1
FF-GAN (fuse)	85.2 ± 1.0	66.3 ± 3.3	90.2 ± 0.6	95.4 ± 0.5

(1) Face Frontalization

Multi-PIE:



LFW:



AFLW2000:



IJB-A:



References:

- [1] T. Hassner, S. Harel, E. Paz, and R. Enbar. "Effective face frontalization in unconstrained images". In *CVPR*, 2015.
- [2] X. Zhu, Z. Lei, J. Yan, D. Yi, and S. Li. "High-Fidelity Pose and Expression Normalization for Face Recognition in the Wild". In *CVPR*, 2015.