

Realistic Face Reconstruction from Facial Embeddings via Diffusion Models

Dong Han^{1,2}, Yong Li^{1*}, Joachim Denzler²

¹Data Protection Technology Lab, Huawei Technologies Düsseldorf, Germany

²Computer Vision Group, Friedrich Schiller University Jena, Germany
{dong.han2, yong.li1}@huawei.com, {dong.han, joachim.denzler}@uni-jena.de

Abstract

With the advancement of face recognition (FR) systems, privacy-preserving face recognition (PPFR) systems have gained popularity for their accurate recognition, enhanced facial privacy protection, and robustness to various attacks. However, there are limited studies to further verify privacy risks by reconstructing realistic high-resolution face images from embeddings of these systems, especially for PPFR. In this work, we propose the face embedding mapping (FEM), a general framework that explores Kolmogorov-Arnold Network (KAN) for conducting the embedding-to-face attack by leveraging pre-trained Identity-Preserving diffusion model against state-of-the-art (SOTA) FR and PPFR systems. Based on extensive experiments, we verify that reconstructed faces can be used for accessing other real-world FR systems. Besides, the proposed method shows the robustness in reconstructing faces from the partial and protected face embeddings. Moreover, FEM can be utilized as a tool for evaluating safety of FR and PPFR systems in terms of privacy leakage. All images used in this work are from public datasets.

Introduction

The advancement of artificial intelligence has brought attention to the security and privacy concerns associated with biometric authentication systems (Laishram et al. 2024; Wang et al. 2024b), specifically for face recognition (FR) (Rezgui, Strisciuglio, and Veldhuis 2024). FR systems generate the template for each identity for comparing different faces and to authenticate query faces. These face templates or face embeddings are considered as one type of biometric data that is frequently produced by *black-box* models (e.g., convolutional neural networks (CNNs) and deep neural networks (DNNs) models). Existing common threats to the face embeddings are sensitive information retrieval attacks (Terhöst et al. 2020) (extract soft-biometric information such as sex, age, race, etc.) or face reconstruction attacks (recover the complete face image from an embedding). To increase the privacy and security level of FR, privacy-preserving face recognition (PPFR) systems (Ji et al. 2022; Mi et al. 2023; Han, Li, and Denzler 2024; Mi et al. 2024) have been proposed. However, most PPFR methods only focus on con-




Figure 1: Sample face images from the CelebA-HQ dataset (first row) and their corresponding reconstructed face images from face embeddings of PPFR model DCTDP. The orange color value indicates confidence score (higher is better) given by commercial API Face++.

cealing visual information from input face images to systems while embeddings are not being protected directly. To protect face embeddings, various protection algorithms are proposed. Hash-based method (Jin et al. 2017) transfers embeddings into discrete index hashed code. Homomorphic Encryption (HE)-based method (Shahreza et al. 2022) encrypts embedding into ciphertext. However, these methods require high computational cost and degrade recognition. Recently, transform-based methods offers more efficient protection. PolyProtect (Hahn and Marcel 2022) is based on multivariate polynomials with user-specific parameters to map embeddings. MLP-Hash (Shahreza, Hahn, and Marcel 2023) utilizes user-specific randomly-weighted multi-layer perceptron (MLP) to rotate and binarize embeddings. SlerpFace (Zhong et al. 2025) employs spherical linear interpolation to rotate embedding towards noise.

Current face reconstruction methods focus on face image reconstruction from embeddings of normal FR models (without designed operations for privacy protection on either input face images or face embeddings). Deconvolutional neural network NbNet (Mai et al. 2018) utilizes the decon-

*Corresponding author.

volution to reconstruct face images from deep templates. End-to-end CNN-based method (Shahreza, Hahn, and Marcel 2022) combines cascaded convolutional layers and deconvolutional layers to improve reconstruction. Moreover, the learning-based method (Shahreza and Marcel 2024a) can reconstruct the underlying face image from a protected embedding that is protected by embedding protection mechanisms (Jin, Ling, and Goh 2004; Shahreza, Hahn, and Marcel 2023; Shahreza et al. 2022). Nevertheless, the reconstructed faces from these methods suffer from noisy and blurry artifacts, which degrade the image naturalness. Generative adversarial network (GAN)-based approach FaceTI (Otroshi Shahreza and Marcel 2024) trains a mapping network to transfer face embedding to the latent space of a pre-trained face generation network. However, they only test their method on normal FR systems and the whole training process is resource-intensive. Recently, MAP2V explores reconstructing face images from PFFR systems by employing GAN model without training. Nevertheless, it requires long inference time which limits its potential in real-time application. Shahreza, George, and Marcel leverage a pre-trained face foundation model (Papantoniou et al. 2024) for face reconstruction. However, it is only tested on normal FR systems. In our work, we choose FaceTI and MAP2V as the main comparisons due to their superiority in realistic face image reconstruction compared with previous CNN-based methods. Considering the above motivations, we propose a face reconstruction framework to generate realistic face identity (ID) based on leaked face embeddings from both FR and PFFR models, by utilizing a pre-trained ID-Preserving diffusion model, IPA-FaceID (Ye et al. 2023). As depicted in Fig. 2, we feed training face images to both IPA-FR (default FR of IPA-FaceID) and target FR models. The initial output face embedding from the target FR model is transferred by the Face Embedding Mapping (FEM) model before performing loss optimization. During the inference stage, the leaked embedding from the target FR or PFFR model can be mapped by trained FEM and directly used by IPA-FaceID to generate realistic face images. We verify the effectiveness of face reconstruction by applying impersonation attacks to real-world FR systems. Besides, reconstructed face images by FEMs can also bypass the commercial face comparison API, Face++¹ as shown in Fig. 1.

Our key contributions are:

- We propose FEM: a face embedding mapping framework that maps the arbitrary embedding of target FR or PFFR for realistic face reconstruction by utilizing pre-trained ID-Preserving diffusion model.
- We exploit the potential of KAN for face embedding mapping and showcase the efficacy of the FEM-KAN model for non-linear mapping compared with SOTA face reconstruction models.
- Through extensive experiments, we demonstrate that proposed method improves on face reconstruction performance and show the generalization to partial and protective embedding

¹<https://www.faceplusplus.com>

Related Work

Identity-Preserving T2I Diffusion Models

Existing text-to-image (T2I) models still have limitations in generating accurate and realistic images due to the limited information expressed by text prompts. IP-Adapter (Ye et al. 2023) proposes decoupled cross-attention to embed image feature to a pre-trained T2I diffusion model by adding a new cross-attention layer. It utilizes a trainable projection model to map the image embedding that extracted by a pre-trained CLIP image encoder (Radford et al. 2021) to a sequence image feature. Later on, IPA-FaceID² is developed for customized face image generation by integrating face information through face embedding extracted from a FR model instead of CLIP image embedding. IPA-FaceID has the ability to produce corresponding diverse styles of image based on a given face image and text prompt. InstantID (Wang et al. 2024a) proposes a trainable lightweight module for transferring face features from the frozen face encoder into the same space of the text token. It leverages modified ControlNet (Zhang, Rao, and Agrawala 2023) for identity preservation. Arc2Face (Papantoniou et al. 2024) is dedicated for ID-to-face generation by using only ID embedding. It fixes the text prompt with a frozen pseudo-prompt “a photo of $\langle id \rangle$ person” where placeholder $\langle id \rangle$ token embedding is replaced by ArcFace embedding of image prompt.

From Deep Face Embeddings to Face images

Extracting images from deep face embeddings is challenging for naive deep learning networks e.g., UNet (Ronneberger, Fischer, and Brox 2015). CNN-based network (Shahreza, Hahn, and Marcel 2022; Shahreza and Marcel 2023) is introduced to reconstruct face images from corresponding embeddings that are extracted from the FR model by end-to-end training. With more restrictions on the embedding leakage of FR models, previous work (Shahreza and Marcel 2024b) attempts to reconstruct the underlying face image from partial leaked face embeddings by using a similar face reconstruction network as (Shahreza, Hahn, and Marcel 2022). However, reconstructed face images from these methods are highly blurred. Furthermore, DSCasConv (cascaded convolutions and skip connections) is proposed to reduce the blurring (Shahreza, Hahn, and Marcel 2024). However, it still has noticeable blurry artifacts around face contour. Most of CNN-based face reconstruction methods focus on reconstructing low-resolution face images. For more realistic face reconstruction, FaceTI (Otroshi Shahreza and Marcel 2024) takes advantage of GAN model to generate face images from the deep face embedding. They employ the pre-trained StyleGAN3 (Karras et al. 2021) network to establish a mapping from facial embeddings to the intermediate latent space of StyleGAN. Recently, MAP2V (Zhang et al. 2024) is a training-free approach that explores reconstructing face images from PFFRs by utilizing StyleGAN to construct prior space and ranks the top-k latent vectors for reconstruction. Shahreza, George, and Marcel train an adapter (a linear layer) to map the face embedding and utilize a pre-trained face foundation model for reconstruction.

²<https://huggingface.co/h94/IP-Adapter-FaceID>

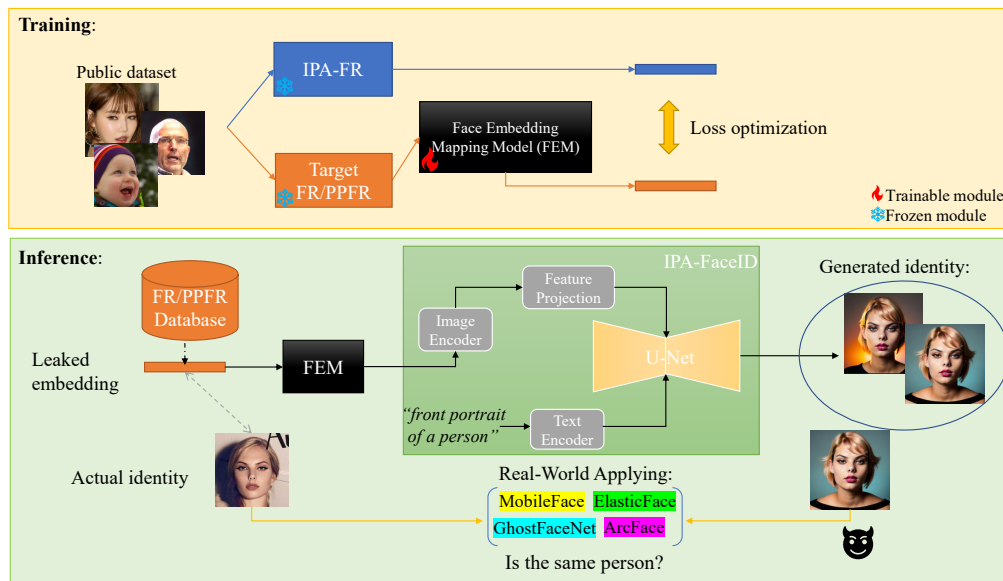


Figure 2: Pipeline of face embedding mapping (FEM). In training, FEM learns to map between target model and the default FR of IPA-FaceID. During inference stage, trained FEM can directly reconstruct complete face images from the leaked embedding.

Approach

Problem Formulation

Attacker’s Goal. Consider a target FR system, the attacker aims to reconstruct a complete face image from a face embedding. Then the reconstructed face image is used to query the same or a different FR system.

Attacker’s Knowledge. The attacker has access to a leaked face embedding of a user registered in the FR system’s database. The attacker only has *black-box* knowledge of the feature extractor model in the same FR system, i.e., can only obtain features from query image. The attacker can inject the reconstructed face image into the target FR systems.

Attacker’s Strategy. Though the attacker has no access to the private face dataset that the target model is being trained on, attacker can collect a public dataset from the Internet or generate a synthetic face dataset for training a face reconstruction model to map face embeddings and reconstruct underlying face images. Then, the attacker can utilize the reconstructed facial images to gain access to the target FR system.

Kolmogorov-Arnold Theorem Preliminaries

The Kolmogorov-Arnold theorem (Liu et al. 2024) states that any continuous function may be expressed as a combination of a finite number of continuous univariate functions. For every continuous function $f(x)$ defined in the n -dimensional real space, where $x = (x_1, x_2, \dots, x_n)$, it can be represented as a combination of a univariate continuous function Φ and a sequence of continuous bivariate functions x_i and $\phi_{q,i}$. The theorem demonstrates the existence of such a representation: $f(x) = \sum_q \Phi_q(\sum_i \phi_{q,i}(x_i))$. It suggests that even sophisticated functions in high-dimensional




Figure 3: Face embedding distributions before and after mapped by FEMs. Visualized by UMAP (McInnes, Healy, and Melville 2018).

spaces can be reconstructed through a sequence of lower-dimensional function operations. The mapping between face embeddings from different systems can be exactly represented as a composition of univariate functions and additions. For face embeddings (typically high-dimensional but structured), this decomposition can better capture complex, non-linear relationships between embedding spaces.

Face Embedding Mapping (FEM)

Face embedding is a vector that represents facial information associated with a corresponding identity and stores soft-biometric information. Previous work (Terhörst et al. 2020) shows demographic information and social traits can be extracted by analyzing face embeddings of FR models. And prediction accuracy is associated with encoding ability of feature extractor. PPFR models employ different types of approaches such as transformation, noise, perturbation to protect privacy. Due to the nature of protection mechanism, i.e., losing information in image space or embedding space,




Figure 4: Two variants of FEM models and the process of embedding-to-embedding mapping. (a) FEM-MLP has fixed activation function. (b) FEM-KAN has learnable activation function at edges to achieve accurate non-linear mapping. (c) The direction of embedding mapping optimized by distance towards to “ground truth” face embedding e_i .

PPFR models sacrifice a certain level of recognition ability. Consequently, soft-biometric information of embeddings can be diminished. Intuitively, it becomes more difficult to extract underlying facial attributes or reconstruct facial image based on PPFR embeddings (Mi et al. 2024; Zhang et al. 2024) than normal FR embeddings. Ideally, embeddings that are extracted from different face images of the same identity should be close and far for those that computed from different ones. Existing SOTA FR and PPFR networks utilize similar structures of backbone to extract features from the face image and compute the face embedding. We assume there is a transformation or mapping algorithm between embeddings from the same identity that are extracted by different backbones. Inspired from (Papantoniou et al. 2024) and (Liu et al. 2024), we propose FEM-MLP and FEM-KAN showing in Fig. 4 to learn the mapping relation of embedding distributions from different FR and PPFR backbones. Then trained FEMs can map face embedding from the initial domain into the corresponding target domain of the pre-trained IPA-FaceID diffusion model in order to generate face images, see in Fig. 3. Depending on the effectiveness of FEMs, the mapped embedding can fall into the target domain and boundary region. The target domain represents mapped embedding can be used for ID-Preserving face image generation that can fool the evaluation FR systems while boundary region indicates mapped embedding is not sufficient for ID-Preserving face image generation but human-like image generation.

For training our face reconstruction framework, considering a target FR or PPFR model $\Gamma'(\cdot)$ and the default FR (IPA-FR) $\Gamma(\cdot)$ model of IPA-FaceID, for training image dataset $\mathcal{I} = I_i$, we can generate the embedding distribution $\mathcal{D}(e_i)$ as well as $\mathcal{D}'(e'_i)$ by extracting face embeddings from all face images in public or synthetic face dataset \mathcal{I} , where e_i and e'_i denote the output face embeddings from $\Gamma(\cdot)$ and $\Gamma'(\cdot)$. In order to enable target $\Gamma'(\cdot)$ model to generate realistic target identity face images by leveraging ID-Preserving capability of the pre-trained IPA-FaceID, the target embedding e'_i extracted from $\Gamma'(\cdot)$ should be mapped close to the corresponding embedding e_i that represents the same face

identity. Therefore, we should minimize the distance between $\hat{\mathcal{D}}(\hat{e}_i) = \mathcal{M}(\mathcal{D}'(e'_i))$ and $\mathcal{D}(e_i)$, where $\mathcal{M}(\cdot)$ and \hat{e}_i denote FEM and mapped face embedding, respectively, by following reconstruction loss function (Mean Square Error):

$$\mathcal{L}_{\text{MSE}}(e_i, \hat{e}_i) = \frac{\sum_{i=0}^{N-1} (e_i - \hat{e}_i)^2}{N} \quad (1)$$

Experiments

Experimental Details

Our goal is to reconstruct complete face image from embedding of target FR and PPFR models. The generated face images are used to fool other real-world FR systems in order to verify the reconstruction performance. We mainly conduct three different experiments to verify the proposed method as follows:

- For evaluating the effectiveness, we reconstruct five face images for each embedding that extracted from target FR and PPFR models and inject the generated faces to the test FR models for performing transferable face verification.
- For evaluating the generalization, we train models on Flickr-Faces-HQ (FFHQ) (Karras, Laine, and Aila 2019) dataset and test on CelebA-HQ (Karras 2017), with 1000 images of never-before-seen identities. Moreover, we also investigate the impact of markup to face reconstruction on LADN dataset (Gu et al. 2019).
- For evaluating the robustness, we conduct experiments to reconstruct faces from partial embeddings and protected embeddings, e.g., PolyProtect (Hahn and Marcel 2022), MLP-Hash (Shahreza, Hahn, and Marcel 2023) and SlerpFace (Zhong et al. 2025). Besides, we also utilize embeddings of facial images safeguarded by the facial privacy algorithm Fawkes (Shan et al. 2020) to reconstruct faces.

Settings

Baselines. We compare FEMs with **FaceTI** (2024) and training-free method **MAP2V** (Zhang et al. 2024). We train FaceTI for 20 epochs with batch size 6, using default configurations in official implementation³. For MAP2V, we use the same setting as in official implementation⁴. To showcase the effectiveness of FEMs, we also directly use pre-trained IPA-FaceID to generate face images for comparison, which is denoted as *None*.

Datasets. We conduct experiments on two face benchmarks, i.e., CelebA-HQ (Karras 2017), LADN (Gu et al. 2019).

Target models. For target models that we aim to reconstruct face image from can be categorized into normal FR models such as **IRSE50** (Hu, Shen, and Sun 2018), **IR152** (Deng et al. 2019) and PPFR models including **DCTDP** (Ji et al. 2022), **HFCF** (Han, Li, and Denzler 2024), **Partial-Face** (Mi et al. 2023) and **MinusFace** (Mi et al. 2024).

³https://gitlab.idiap.ch/bob/bob.paper.neurips2023_face_ti

⁴<https://github.com/Beauty9882/MAP2V>

Evaluation metrics. Following (Shamshad, Naseer, and Nandakumar 2023; Sun et al. 2024), we employ the attack success rate (ASR) as a metric to evaluate attack efficacy of reconstructed face images from target FR and PFR systems. ASR is defined as the fraction of generated faces successfully recognized by real-word FR systems. We generate five images for each embedding. When determining the ASR, we establish a False Acceptance Rate (FAR) of 0.01 for each FR model. We have selected four widely-used public FR models as real-word models for face verification to measure reconstruction performance. These models are ElasticFace (**EF**) (Boutros et al. 2022), MobileFace (**MF**) (Chen et al. 2018), GhostFaceNet (**GF**) (Alansari et al. 2023), ArcFace (**AF**) (Deng et al. 2019). We use implementation from *DeepFace*⁵ for last two models.

Implementation details. FR *buffalo_l*⁶ is selected as the default FR model of IPA-FaceID. We choose *faceid_sd15*⁷ checkpoint for IPA-FaceID which takes face embedding and text as input. In order to effectively generate face in proper angle, we fix the text prompt as “front portrait of a person” for all the experiments. As our finding, few layers are sufficient for FEMs achieving good mapping performance. Therefore, we implement FEM-KAN with three KAN layers (follow *efficient_kan*⁸) and FEM-MLP with three MLP layers. We use the GELU activation function and add 1D batch normalization to FEM-MLP. We train two variant FEM models with 90% FFHQ dataset for 20 epochs. For optimizers, we use AdamW with initial learning rate 10^{-2} and the exponential learning rate decay is set to 0.8. Our experiments are conducted on a Tesla V100 GPU with 32G memory using PyTorch framework, setting the batch size to 128 for FEMs.

Results

Performance on Privacy Attacks

As shown in Tab. 1, our proposed FEMs can effectively reconstruct face images from target FR models and outperform FaceTI and MAP2V. Besides, FEMs also maintain high ASRs on various target PFR models. It shows that even images transferred into frequency domain, the corresponding face embedding still contains information can be used for high-quality face reconstruction. FEMs achieve higher average ASR on all PFR models than existing SOTA method MAP2V, especially on HFCF and MinusFace. We exclude training PFR models with FaceTI due to the constraints of our computational resources and the extremely long training time required by this method, details about resource requirements are in the ablation study. We show a visual comparison of reconstructed face images in Fig. 5.

Makeup Reconstruction From Face Embeddings

Previous work shows (Chen et al. 2017) makeup can be used as markup presentation attacks, e.g., for impersonation. For

⁵<https://github.com/serengil/deepface>

⁶<https://github.com/deepinsight/insightface>

⁷<https://huggingface.co/h94/IP-Adapter-FaceID/tree/main>

⁸<https://github.com/Blealtan/efficient-kan>

Target FR	Method	MF	EF	GF	AF	Average
IRSE50	None	10.7	7.3	3.0	9.1	7.5
	FaceTI	93.4	80.8	49.6	66.8	72.7
	MAP2V	94.0	86.2	59.3	72.0	77.9
	FEM-MLP	98.0	91.8	62.6	73.4	81.5
	FEM-KAN	99.2	93.8	65.7	76.1	83.7
IR152	None	9.2	5.8	7.6	2.3	6.2
	FaceTI	85.2	74.0	43.4	61.2	66.0
	MAP2V	92.3	82.5	53.0	67.2	73.8
	FEM-MLP	94.1	78.1	48.6	64.6	71.4
	FEM-KAN	95.0	85.2	58.4	70.3	77.2
Target PFR						
DCTDP	None	6.8	5.4	2.9	7.1	5.6
	MAP2V	94.0	85.5	59.4	74.3	78.3
	FEM-MLP	97.3	91.8	67.9	77.7	83.7
	FEM-KAN	98.5	91.8	68.4	78.8	84.4
	HFCF	None	7.5	4.2	5.1	11.4
MAP2V		76.3	15.4	5.3	14.8	28.0
FEM-MLP		97.0	89.2	61.8	74.2	80.6
FEM-KAN		98.3	90.7	66.5	76.9	83.1
PartialFace		None	7.9	5.8	1.8	6.6
	MAP2V	93.1	84.3	60.0	71.4	77.2
	FEM-MLP	98.7	88.5	61.5	71.6	80.1
	FEM-KAN	99.5	90.5	68.0	77.6	83.9
	MinusFace	None	3.5	5.1	2.1	10.3
MAP2V		68.0	4.8	2.3	5.6	20.2
FEM-MLP		94.4	68.5	45.9	59.4	67.1
FEM-KAN		96.5	71.3	44.5	58.1	67.6

Table 1: Evaluations of ASR for black-box attacks to FR and PFR models on CelebA-HQ dataset.

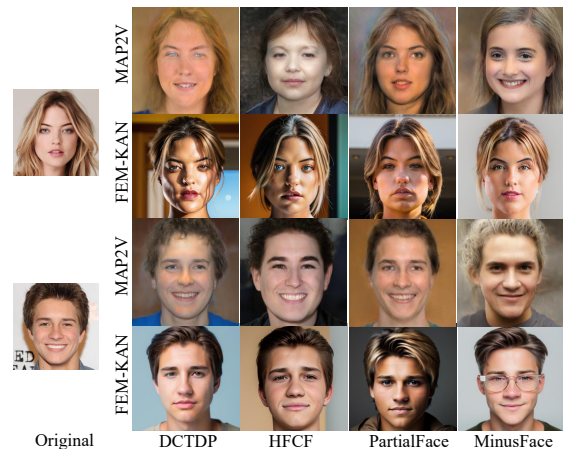


Figure 5: Visual Comparison of reconstructed faces from CelebA-HQ dataset.

instance, attacker can apply specific markup such that the face of attacker looks similar to that of a target subject. However, there is no existing work studying how markup will affect the face reconstruction. Therefore, apart from reconstructing natural-looking face images from face embeddings, we also consider a more challenging scenario to evaluate whether the proposed method can also reconstruct markups applied to face images that were used for extracting face embeddings. As shown in Tab. 2, we demonstrate

Dataset	Method	MF	EF	GF	AF	Average
LADN-NM	FaceTI	98.5	75.4	50.6	73.4	74.5
	MAP2V	98.2	93.1	66.8	79.0	84.3
	FEM-MLP	100.0	94.0	72.5	83.5	87.5
	FEM-KAN	100.0	94.9	78.7	92.5	91.5
LADN-M	FaceTI	91.0	57.7	30.1	46.8	56.4
	MAP2V	97.2	85.6	51.3	67.9	75.5
	FEM-MLP	99.4	87.9	63.7	73.5	81.1
	FEM-KAN	100.0	94.9	69.3	76.3	85.1

Table 2: ASR performance on reconstructing face images with markup. **LADN-NM** and **LADN-M** denote non-markup and markup datasets.




Figure 6: (a) Reconstructed face comparison with different embedding leakages. (b) ASR performance evaluated by ArcFace. IRSE50 is target model.

the impact of markup to face reconstruction. Markup negatively affects the ASR of FaceTI resulting in significant accuracy decreasing, i.e., 20.5% on GF and 26.6% on AF, respectively. Moreover, FaceTI only achieves 56.4% average ASR on markup dataset. MAP2V also suffers a noticeable average ASR drop around 8.8%. In contrast, our proposed FEMs are more robust to markup and maintain high performance in reconstructing face images with markup, e.g., 85.1% ASR in average.

Face Reconstruction from Partial Leaked Embeddings

Previous experiments are based on the assumption that the adversary can gain access to the complete face embeddings. Nevertheless, in some real-world scenarios, a complete face embedding is difficult to acquire, but rather to access a portion of the embedding. For example, face embeddings of the FR system are split and stored on different servers for data protection like the situation considered in (Shahreza and Marcel 2024b). We assume the adversary already trained FEMs on complete embeddings of the target FR or PPF model. To further test non-linear mapping ability and face reconstruction, we only use partial leaked embeddings (e.g., discarding the half values in an embedding vector in the case of 50% leakage) as input to trained FEMs. To match the input shape of FEMs, we append zeros to the end of each leaked embedding vector. In Fig. 6 (a), the reconstructed face images can reveal privacy-sensitive information about the underlying user, such as gender, race, etc. FaceTI and MAP2V can reconstruct general attributes (gender, age and facial expression) in the case of high leakage, e.g., 70% to 90% while our method is able to reconstruct more detailed

Protection	Method	MF	EF	GF	AF
PolyProtect	FaceTI	45.6	3.3	0.8	4.1
	MAP2V	28.6	4.4	3.6	4.3
	FEM-MLP	48.3	4.7	6.8	15.6
	FEM-KAN	50.3	7.1	5.6	15.4
MLP-Hash	FaceTI	66.8	0.7	0.1	0.4
	MAP2V	48.1	0.6	0.3	1.5
	FEM-MLP	80.4	53.0	51.9	64.5
	FEM-KAN	82.1	54.7	56.5	71.6
SlerpFace	FaceTI	24.6	1.2	0.3	0.1
	MAP2V	11.4	0.0	0.1	0.1
	FEM-MLP	78.6	7.1	4.5	14.1
	FEM-KAN	79.4	9.3	7.8	15.4

Table 3: ASR performance on protected face embeddings.

attributes similar to the original identity. Under 50% leakage scenario, FaceTI start to generate unified and normalized faces for embeddings belonging to different identities. In contrast, FEM can still reconstruct the general facial structures of identity. However, the generated face tends to have noticeable artifacts when leakage is lower than 30%. Fig. 6 (b) reports ASR to evaluate the incomplete leaked embedding mapping ability of FEMs. With an increased percentage of embedding leakage, the number of generated face images that can fool the test FR is reduced. FaceTI suffers the significant impact with decreased embedding leakage, i.e., 22% ASR reduction when 50% leakage compared with MAP2V and FEMs. For the extreme low leakage below 30%, FEMs still achieve relatively high ASR (32.5% for FEM-KAN). Therefore, our method is more robust than GAN-based method and able to reconstruct facial information from incomplete embeddings.

Face Reconstruction from Protected Embeddings

Considering more strict access to the original embeddings that were directly computed by the feature extractor of PPF models, we evaluate FEMs on face embeddings that being protected by particular embedding protection algorithms such as **PolyProtect** (Hahn and Marcel 2022), **MLP-Hash** (Shahreza, Hahn, and Marcel 2023) and **SlerpFace** (Zhong et al. 2025). We train FEMs directly on protected face embeddings derived from these protection algorithms based on PPF HFCF model. For PolyProtect, we generate the user-specific pair for each identity in the testing dataset. After mapping the original face embedding from PPF model, the protected embedding from PolyProtect has reduced dimension, 508 in our setting. During training, we append other four zeros to the end of the protected embedding to maintain the length of vector. For MLP-Hash method, we set one-hidden layer with 512 neurons and fix the seed for all identities. As for SlerpFace, we use the default settings in original implementation. Tab. 3 reports the ASR on the protected embeddings. It is worth noticing that FEMs still can achieve high face reconstruction against MLP-Hash and have comparable ASR with the ones on unprotected embeddings in Tab. 1. Moreover, FEMs have the highest ASRs against three protection algorithms than FaceTI and MAP2V.

Protection	Method	MF	EF	GF	AF
Fawkes	FaceTI	85.0	16.3	6.3	13.0
	MAP2V	96.3	44.9	17.4	25.7
	FEM-MLP	97.0	42.0	16.7	24.2
	FEM-KAN	97.0	47.1	19.0	27.9

Table 4: ASR performance on protected face images.

	Training Time	Memory	Inference Time
FaceTI	51 hrs	25383 MiB	3.2s
MAP2V	0	32133 MiB	111s
FEMs	3 hrs	4325 MiB	2.6s

Table 5: Training resource requirements when training (single epoch) on 90% FFHQ dataset and inference time on reconstruct a single image.

Face Reconstruction from Protected Facial Images

Unauthorized FR systems identifying and recognizing online photos pose a threat to personal security and privacy in the digital world. Most of existing countermeasures for protecting privacy of facial images are based on introducing perturbations (Madry et al. 2018; Shan et al. 2020) or adding markups (Sun et al. 2024) in pixel level to misguide unauthorized FR systems. Intuitively, it is more difficult to reconstruct facial image from leaked embedding of protected image. In this section, we focus on evaluating FEMs against SOTA facial protection algorithm **Fawkes** (Shan et al. 2020) on IRSE50. It is a *black-box* cloaking algorithm to add imperceptible perturbations. Considering computational cost and strength of privacy protection, we apply Fawkes to CelebA-HQ using mid level perturbations to generate cloaked version of dataset. As depicted in Tab. 4, ASRs are considerably reduced across on different public FR systems compared with performance (see in Tab. 1) on reconstructing face images from embeddings of corresponding clean dataset. It is worth to notice that reconstructed face images still maintain high ASRs to access MobileFace, which can threaten to other lightweight FR systems that designed for running face recognition algorithms with limited hardware resources, e.g., phone and tablet. Therefore, the more powerful face image protection algorithms should be developed to further reduce the risk of face reconstruction attack without introducing noticeable artifacts.

Ablation Study

Training and inference efficacy. To produce human-like images using pre-trained StyleGAN3 model, FaceTI requires training an additional critic network (Wasserstein GAN (Arjovsky, Chintala, and Bottou 2017)) together with the mapping network simultaneously, which demands high computational cost and is resource-intensive. In contrast, FEMs require **17** \times less training time and **5.8** \times less GPU memory compared with FaceTI, see in Tab. 5. We also measure inference time that reconstructing a complete face from an embedding. MAP2V takes extreme long time for single image generation. Our method can reconstruct one image within 3 seconds, which is around **42** \times faster than MAP2V.

Method	MF	EF	GF	AF	Average
FaceTI	66.4	54.4	43.6	69.8	58.6
MAP2V	86.8	66.2	52.8	54.2	65.0
FEM-MLP	89.0	67.2	38.4	46.6	60.3
FEM-KAN	88.0	69.0	45.4	50.8	63.3

Table 6: Evaluations of ASR for black-box attacks to IRSE50 on LFW dataset.




Figure 7: The percentage of reconstructed images passing face anti-spoofing system.

Reconstruction from low-resolution facial images. For scenario that leaked embedding is from low-resolution face, we randomly select 500 images from LFW (Huang et al. 2008), with image size in 112×112 . As shown in Tab. 6, we observe dropped ASRs compared with the performance on high-resolution CelebA-HQ images. **Face anti-spoofing.** Face Anti-Spoofing (FAS) is a critical security measure designed to detect and prevent attempts to deceive FR systems. By assuming real-world FR systems incorporate FAS to reject potential fake face images before recognition, we conduct experiment with FASNet (Lucena et al. 2017) to evaluate the attacking effectiveness by passing ratio, which is defined as the fraction reconstructed faces bypassing the FAS system. In Fig. 7, we also add CNN-based method (Shahreza and Marcel 2023) to demonstrate why it is not practical for real-world attack. Since reconstructed face images from CNN-based method are not realistic (e.g., blurry, noisy artifacts) and have low resolutions, these images are easily rejected by FAS system. MAP2V faces the same issue when reconstructing from LFW dataset.

Conclusion

In this paper, we propose a new framework called FEM to reconstruct high-resolution, realistic face images from face embeddings in both FR and PFR systems. FEM can project arbitrary face embeddings for face reconstruction by leveraging a pre-trained face image generation network, IPA-FaceID. We measure the effectiveness in different scenarios, including *black-box* embedding-to-face attacks, out-of-distribution generalization, reconstructing from partial leaked embeddings, protected embeddings and protected facial images. Extensive experiments demonstrate that FEM outperforms SOTAs in both FR and PFR scenarios. Moreover, FEM framework can be served as an evaluation tool to verify the robustness of existing PFRs and privacy protection algorithms.

References

- Alansari, M.; Hay, O. A.; Javed, S.; Shoufan, A.; Zweiri, Y.; and Werghi, N. 2023. Ghostfacenets: Lightweight face recognition model from cheap operations. *IEEE Access*, 11: 35429–35446.
- Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein generative adversarial networks. In *International conference on machine learning*, 214–223. PMLR.
- Boutros, F.; Damer, N.; Kirchbuchner, F.; and Kuijper, A. 2022. Elasticface: Elastic margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1578–1587.
- Chen, C.; Dantcheva, A.; Swearingen, T.; and Ross, A. 2017. Spoofing faces using makeup: An investigative study. In *2017 IEEE International Conference on Identity, Security and Behavior Analysis (ISBA)*, 1–8. IEEE.
- Chen, S.; Liu, Y.; Gao, X.; and Han, Z. 2018. Mobilefacenets: Efficient cnns for accurate real-time face verification on mobile devices. In *Biometric Recognition: 13th Chinese Conference, CCBR 2018, Urumqi, China, August 11-12, 2018, Proceedings 13*, 428–438. Springer.
- Deng, J.; Guo, J.; Xue, N.; and Zafeiriou, S. 2019. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4690–4699.
- Gu, Q.; Wang, G.; Chiu, M. T.; Tai, Y.-W.; and Tang, C.-K. 2019. Ladm: Local adversarial disentangling network for facial makeup and de-makeup. In *Proceedings of the IEEE/CVF International conference on computer vision*, 10481–10490.
- Hahn, V. K.; and Marcel, S. 2022. Towards protecting face embeddings in mobile face verification scenarios. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 4(1): 117–134.
- Han, D.; Li, Y.; and Denzler, J. 2024. Privacy-Preserving Face Recognition in Hybrid Frequency-Color Domain. In *International Conference on Computer Vision Theory and Applications (VISAPP)*, 536–546. INSTICC, SciTePress. ISBN 978-989-758-679-8.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-Excitation Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Huang, G. B.; Mattar, M.; Berg, T.; and Learned-Miller, E. 2008. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In *Workshop on faces in 'Real-Life' Images: detection, alignment, and recognition*.
- Ji, J.; Wang, H.; Huang, Y.; Wu, J.; Xu, X.; Ding, S.; Zhang, S.; Cao, L.; and Ji, R. 2022. Privacy-preserving face recognition with learnable privacy budgets in frequency domain. In *European Conference on Computer Vision*, 475–491. Springer.
- Jin, A. T. B.; Ling, D. N. C.; and Goh, A. 2004. Biohashing: two factor authentication featuring fingerprint data and tokenised random number. *Pattern recognition*, 37(11): 2245–2255.
- Jin, Z.; Hwang, J. Y.; Lai, Y.-L.; Kim, S.; and Teoh, A. B. J. 2017. Ranking-based locality sensitive hashing-enabled cancelable biometrics: Index-of-max hashing. *IEEE Transactions on Information Forensics and Security*, 13(2): 393–407.
- Karras, T. 2017. Progressive Growing of GANs for Improved Quality, Stability, and Variation. *arXiv preprint arXiv:1710.10196*.
- Karras, T.; Aittala, M.; Laine, S.; Härkönen, E.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2021. Alias-free generative adversarial networks. *Advances in neural information processing systems*, 34: 852–863.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4401–4410.
- Laishram, L.; Shaheryar, M.; Lee, J. T.; and Jung, S. K. 2024. Toward a Privacy-Preserving Face Recognition System: A Survey of Leakages and Solutions. *ACM Computing Surveys*.
- Liu, Z.; Wang, Y.; Vaidya, S.; Ruehle, F.; Halverson, J.; Soljačić, M.; Hou, T. Y.; and Tegmark, M. 2024. Kan: Kolmogorov-arnold networks. *arXiv preprint arXiv:2404.19756*.
- Lucena, O.; Junior, A.; Moia, V.; Souza, R.; Valle, E.; and Lotufo, R. 2017. Transfer learning using convolutional neural networks for face anti-spoofing. In *Image Analysis and Recognition: 14th International Conference, ICIAR 2017, Montreal, QC, Canada, July 5–7, 2017, Proceedings 14*, 27–34. Springer.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *International Conference on Learning Representations*.
- Mai, G.; Cao, K.; Yuen, P. C.; and Jain, A. K. 2018. On the reconstruction of face images from deep face templates. *IEEE transactions on pattern analysis and machine intelligence*, 41(5): 1188–1202.
- McInnes, L.; Healy, J.; and Melville, J. 2018. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.
- Mi, Y.; Huang, Y.; Ji, J.; Zhao, M.; Wu, J.; Xu, X.; Ding, S.; and Zhou, S. 2023. Privacy-preserving face recognition using random frequency components. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19673–19684.
- Mi, Y.; Zhong, Z.; Huang, Y.; Ji, J.; Xu, J.; Wang, J.; Wang, S.; Ding, S.; and Zhou, S. 2024. Privacy-preserving face recognition using trainable feature subtraction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 297–307.
- Otroshi Shahreza, H.; and Marcel, S. 2024. Face reconstruction from facial templates by learning latent space of a generator network. *Advances in Neural Information Processing Systems*, 36.

- Papantoniou, F. P.; Lattas, A.; Moschoglou, S.; Deng, J.; Kainz, B.; and Zafeiriou, S. 2024. Arc2face: A foundation model of human faces. *arXiv preprint arXiv:2403.11641*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Rezgui, Z.; Strisciuglio, N.; and Veldhuis, R. 2024. Enhancing soft biometric face template privacy with mutual information-based image attacks. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1141–1149.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, 234–241. Springer.
- Shahreza, H. O.; George, A.; and Marcel, S. 2025. Face Reconstruction from Face Embeddings using Adapter to a Face Foundation Model. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 5584–5593.
- Shahreza, H. O.; Hahn, V. K.; and Marcel, S. 2022. Face reconstruction from deep facial embeddings using a convolutional neural network. In *2022 IEEE International Conference on Image Processing (ICIP)*, 1211–1215. IEEE.
- Shahreza, H. O.; Hahn, V. K.; and Marcel, S. 2023. Mlp-hash: Protecting face templates via hashing of randomized multi-layer perceptron. In *2023 31st European Signal Processing Conference (EUSIPCO)*, 605–609. IEEE.
- Shahreza, H. O.; Hahn, V. K.; and Marcel, S. 2024. Vulnerability of State-of-the-Art Face Recognition Models to Template Inversion Attack. *IEEE Transactions on Information Forensics and Security*.
- Shahreza, H. O.; and Marcel, S. 2023. Blackbox face reconstruction from deep facial embeddings using a different face recognition model. In *2023 IEEE International Conference on Image Processing (ICIP)*, 2435–2439. IEEE.
- Shahreza, H. O.; and Marcel, S. 2024a. Breaking Template Protection: Reconstruction of Face Images from Protected Facial Templates. In *2024 IEEE 18th International Conference on Automatic Face and Gesture Recognition (FG)*, 1–7. IEEE.
- Shahreza, H. O.; and Marcel, S. 2024b. Face reconstruction from partially leaked facial embeddings. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4930–4934. IEEE.
- Shahreza, H. O.; Rathgeb, C.; Osorio-Roig, D.; Hahn, V. K.; Marcel, S.; and Busch, C. 2022. Hybrid protection of biometric templates by combining homomorphic encryption and cancelable biometrics. In *2022 IEEE International Joint Conference on Biometrics (IJCB)*, 1–10. IEEE.
- Shamshad, F.; Naseer, M.; and Nandakumar, K. 2023. Clip2protect: Protecting facial privacy using text-guided makeup via adversarial latent search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20595–20605.
- Shan, S.; Wenger, E.; Zhang, J.; Li, H.; Zheng, H.; and Zhao, B. Y. 2020. Fawkes: Protecting privacy against unauthorized deep learning models. In *29th USENIX security symposium (USENIX Security 20)*, 1589–1604.
- Sun, Y.; Yu, L.; Xie, H.; Li, J.; and Zhang, Y. 2024. Dif-fAM: Diffusion-based Adversarial Makeup Transfer for Facial Privacy Protection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24584–24594.
- Terhörst, P.; Fährmann, D.; Damer, N.; Kirchbuchner, F.; and Kuijper, A. 2020. Beyond identity: What information is stored in biometric face templates? In *2020 IEEE international joint conference on biometrics (IJCB)*, 1–10. IEEE.
- Wang, Q.; Bai, X.; Wang, H.; Qin, Z.; and Chen, A. 2024a. Instantid: Zero-shot identity-preserving generation in seconds. *arXiv preprint arXiv:2401.07519*.
- Wang, T.; Zhang, Y.; Xiao, X.; Yuan, L.; Xia, Z.; and Weng, J. 2024b. Make Privacy Renewable! Generating Privacy-Preserving Faces Supporting Cancelable Biometric Recognition. In *ACM Multimedia 2024*.
- Ye, H.; Zhang, J.; Liu, S.; Han, X.; and Yang, W. 2023. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*.
- Zhang, H.; Dong, X.; Lai, Y.; Zhou, Y.; Zhang, X.; Lv, X.; Jin, Z.; and Li, X. 2024. Validating Privacy-Preserving Face Recognition under a Minimum Assumption. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12205–12214.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3836–3847.
- Zhong, Z.; Mi, Y.; Huang, Y.; Xu, J.; Mu, G.; Ding, S.; Zhang, J.; Guo, R.; Wu, Y.; and Zhou, S. 2025. Slerpface: Face template protection via spherical linear interpolation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 10698–10706.