PCE-Palm: Palm Crease Energy Based Two-Stage Realistic Pseudo-Palmprint Generation

Jianlong Jin[1, 2*], Lei Shen[2*], Ruixin Zhang[2], Chenglong Zhao[2], Ge Jin[2], Jingyun Zhang[2], Shouhong Ding[2], Yang Zhao[1†], Wei Jia[1†]

[1]Hefei University of Technology, China
[2]Youtu Lab, Tencent
jianlong@mail.hfut.edu.cn, shenlei1996@gmail.com,
{ruixinzhang, lornezhao, maxwelljin, naskyzhang, ericshding}@tencent.com, {yzhao, jiawei}@hfut.edu.cn

Abstract

The lack of large-scale data seriously hinders the development of palmprint recognition. Recent approaches address this issue by generating large-scale realistic pseudo palmprints from Bézier curves. However, the significant difference between Bézier curves and real palmprints limits their effectiveness. In this paper, we divide the Bézier-Real difference into creases and texture differences, thus reducing the generation difficulty. We introduce a new palm crease energy (PCE) domain as a bridge from Bézier curves to real palmprints and propose a two-stage generation model. The first stage generates PCE images (realistic creases) from Bézier curves, and the second stage outputs realistic palmprints (realistic texture) with PCE images as input. In addition, we also design a lightweight plug-and-play line feature enhancement block to facilitate domain transfer and improve recognition performance. Extensive experimental results demonstrate that the proposed method surpasses state-of-the-art methods. Under extremely few data settings like 40 IDs (only 2.5%of the total training set), our model achieves a 29% improvement over RPG-Palm and outperforms ArcFace with 100% training set by more than 6% in terms of TAR@FAR=1e-6.

Introduction

As a promising biometric technology, palmprint recognition has demonstrated broad application prospects (Fei et al. 2018). Big companies like Amazon and Tencent have released palmprint based products. In recent years, deep learning has become the mainstream palmprint recognition technology (Zhong and Zhu 2019). However, a significant bottleneck in the research and application of deep learning-based palmprint recognition lies in the scarcity of large-scale palmprint datasets. Collecting large-scale palmprint datasets may pose a risk of violating user privacy.

To address this challenge, BézierPalm (Zhao et al. 2022) uses parameterized Bézier curves to simulate pseudo-palmprints and pretrain recognition model. By adjusting the control points of Bézier curves according to geometric rules, BézierPalm can generate massive new identities without training data. However, as shown in Fig.1, Bézier curves are noticeably different from real palmprints in terms of both crease distribution and texture. As a result, BézierPalm still requires a non-neglectable amount of data for finetuning.

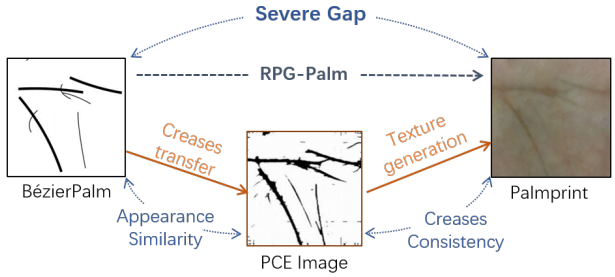

Figure 1: BézierPalm (Zhao et al. 2022), RPG-Palm (Shen et al. 2023), and the proposed two-stage PCE-Palm.

Inspired by BézierPalm, RPG-Palm (Shen et al. 2023) proposes a generation model that directly generates realistic pseudo-palmprints from the Bézier curves with unpaired training. To address this challenge, the generation model incorporates a palmprint recognition model to ensure identity consistency and a conditional modulation generator to improve the intraclass diversity. As a result, RPG-Palm can generate massive realistic palmprints and significantly improve recognition performance. Unfortunately, RPG-Palm still suffers from the severe domain gap and difficulty in unpaired training. Both identity consistency and generation models cannot work well with limited training data.

In this paper, we propose a novel two-stage method to well address the severe domain gap between Bézier curves and realistic palmprints.

Our method introduces a palm crease energy (PCE) domain as an intermediate bridge connecting Bézier curves and realistic palmprints. Therefore, the generation process can be divided into two stages: Bézier curves to PCE images and PCE images to palmprints, as illustrated in Fig.2.

In the first stage, our generation model produces PCE images by transferring the style from Bézier curves to realistic palm creases. Both the PCE images and Bézier curves are binary line-based images, simplifying the task to focus solely on transferring the style from the curves to realistic creases. As depicted in Fig.2, although the task remains unpaired,

*Equal contribution.
†Corresponding authors.

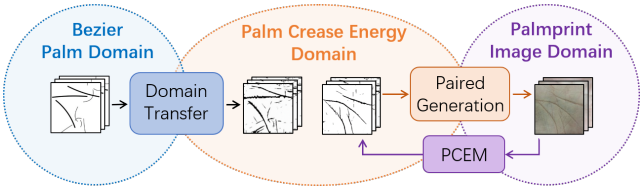

Figure 2: The proposed two-stage generation model introduces the PCE domain to bridge the domain gap.

the difficulty is reduced compared to directly generating realistic creases and skin textures from Bézier curves. In the second stage, we introduce a cycle conditional generation model to transform images from the PCE domain to the realistic palmprint domain. Since PCE images are extracted from real palmprints and share the same distribution as real palm creases, the generation model in the second stage primarily focuses on generating realistic skin textures based on the provided PCE creases. As shown in Fig.2, the paired supervision further reduces the learning difficulty during the process. By dividing the challenging unpaired generation process from Bézier curves to real palmprints into two more tractable stages, the approach enhances the generation performance and reduces the dependence on real data.

Inspired by the modified finite Radon transform (MFRAT) (Jia, Huang, and Zhang 2008), we propose a differentiable PCE extraction module (PCEM) with improved Gaussian-MFRAT kernels to establish a suitable PCE domain. Cooperating with a cycle loss, the PCEM also serves as a method for ensuring identity consistency in the second stage, effectively circumventing the challenge of training a reliable ID discriminator when facing a scarcity of samples. Based on the same idea of PCEM, we also introduce a lightweight line feature enhancement block (LFEB) that can be integrated into deep networks to enhance the perception of line features. The main contributions of this paper are summarized as follows:

- We propose a novel two-stage realistic pseudo-palmprint generation paradigm. To our best knowledge, we are the first to divide the palmprint generation task into two stages: realistic creases generation and realistic textures generation. By introducing the PCE domain as a bridge between Bézier curves and real palmprints, the entire generation process becomes easier and can perform well even with extremely little training data.
- We propose a differentiable PCE module (PCEM) with Gaussian-MFRAT kernels, which can conveniently map real palmprint images to the PCE domain. We also propose a plug-and-play Line Features Enhancement Block (LFEB), which can boost the performance of both generations and recognition models.
- We improve the Bézier curves generation mechanism using prior knowledge of human dermatoglyphics. The distribution of improved curves is closer to real creases.
- Experimental results show that the proposed method surpasses existing SOTA methods under both full sample training and few sample training settings. Under extremely few sample settings like 40 IDs (2.5% of the training set), PCE-Palm outperforms RPG-Palm by 29% and 100%-real-data trained ArcFace by more than 6% (TAR@FAR=1e-6).

Related Work

Palmprint Recognition Methods

Traditional palmprint recognition methods can be roughly divided into two categories, i.e., the local-based methods (Fei et al. 2016; Guo et al. 2009; Kong and Zhang 2004; Jia, Huang, and Zhang 2008; Luo et al. 2016; Sun et al. 2005; Zheng, Kumar, and Pan 2015) that extract effective local features for recognition by manually designing, and the global-based methods (Almeida, Shmarko, and Lomas 2022; Feng et al. 2006; Hu, Feng, and Zhou 2007; Lu, Zhang, and Wang 2003; Sang, Yuan, and Zhang 2009; Wang and Ruan 2006) obtain low-dimensional features from the entire palmprint ROI. Deep Learning based models (Dian and Dongmei 2016; Genovese et al. 2019; Jia et al. 2022; Svoboda, Masci, and Bronstein 2016) train neural networks to extract features of palmprints with improved classification or pair-wise loss (Deng et al. 2019; Shen et al. 2022; Zhong and Zhu 2019). However, further research on palmprint recognition is hindered by the lack of large-scale public datasets.

Data Generation for Recognition Task

With the development of generation methods (Goodfellow et al. 2020; Karras, Laine, and Aila 2019), some models have been applied for producing pseudo samples in the field of biometrics, such as recognition-oriented face and 3D face generation (Deng et al. 2020; Nguyen-Phuoc et al. 2019; Qiu et al. 2021; Yin et al. 2017; Fu et al. 2019, 2021; Geng, Cao, and Tulyakov 2019), and fingerprint synthesis (Bahmani et al. 2021; Engelsma, Grosz, and Jain 2022; Wyzykowski, Segundo, and de Paula Lemes 2021). For palmprint recognition, BézierPalm (Zhao et al. 2022) first synthesizes fake palm creases by using parameterized Bézier curves. RPG-Palm (Shen et al. 2023) further generates realistic pseudo-palmprints by using Bézier curves as the identity control condition. However, due to the severe gap between Bézier curves and real palmprints, RPG-Palm suffers from significant performance degradation in cases with few samples.

Few-shot Generation

With limited data, GAN-based models often face challenges with discriminator overfitting and the imbalance between discrete data spaces and continuous latent distributions (Li et al. 2022; Yang and Wang 2023). These issues result in reduced fidelity and unstable training processes. Recent few-shot improvement methods can be divided into the following three categories, i.e., data augmentation (Dai, Hang, and Guo 2022; Jeong and Shin 2021; Karras et al. 2020; Zhao et al. 2020), regularization (Kong et al. 2022; Yang et al. 2021; Zhang et al. 2019a; Zhao et al. 2021), and transfer learning (Mo, Cho, and Shin 2020; Liu et al. 2019; Ojha et al. 2021). Nevertheless, among prior optimization techniques for few-shot generation, there is still a lack of identity-controllable palmprint generation methods.

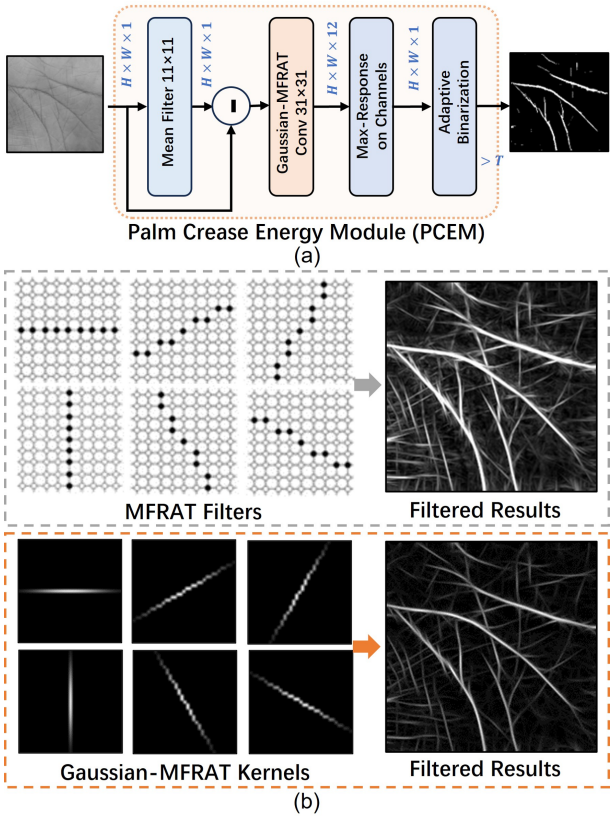

Figure 3: Proposed PCE extraction module (PCEM), (a) the structure of the PCEM, (b) examples of traditional MFRAT filters (Jia, Huang, and Zhang 2008), the proposed Gaussian-MFRAT kernels, and filtered results.

Method

Overall Framework

The framework of the proposed method is illustrated in Fig.2, which is a two-stage model using PCE as an intermediate domain. In the following, we will introduce the PCE domain in detail, and then introduce the structures of generation models in two stages. Finally, the improved Bézier curves generation method is also discussed.

The Construction of Palm Crease Energy Domain

In the PCE domain, we aim to extract binary creases that are similar to Bézier curves from real palmprint images. Inspired by the traditional line extraction algorithm MFRAT (Jia, Huang, and Zhang 2008), which uses a set of manually designed linear filters and employs a maximum response strategy to compute the line orientation energy, we redesign this traditional algorithm into a differentiable PCE extraction module (PCEM) that can be easily embedded into deep networks, as shown in Fig.3 (a).

For an input grayscale image Y (256×256), a 11×11 mean filter is firstly applied to obtain the high-frequency components by subtracting the filtered result from the original Y. Then, the line features are obtained through a line-like convolution layer that consists of several Gaussian-MFRAT kernels along different directions. A max-response operation on channels is presented to achieve non-linearity instead of a common activation layer, which selects the maximum values for each pixel along the channel direction and reproduces a single-channel output that reflects the line orientation energy. Finally, an adaptive binarization is used to obtain the final PCE image. To highlight the principle lines, the binarization threshold T is set according to the top 10% of the values on the whole image. Note that the images generated by PCEM have a black background. To keep with BéizerPalm, we perform color inversion during subsequent training.

Fig.3 (b) illustrates the differences between MFRAT filters and the proposed Gaussian-MFRAT kernels. Traditional MFRAT uses line-like filters with constant values, which are sensitive to noise or tiny variations. The proposed Gaussian-MFRAT kernels are calculated as follows,

$$f(x,y) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(\frac{\| \left((x-x_0), (y-y_0) \right) \|_\infty^2}{2\sigma^2} \right), \quad (1)$$

where $(x,y) \in L(\theta)$ represents the coordinates on the kernel, (x_0, y_0) denotes the central point of it, $L(\theta)$ denotes a line defined on a two-dimensional image plane with an angle of θ, and σ is a hyperparameter. Additionally, $f(x,y) = 0$ when $(x,y) \notin L(\theta)$. Referring to (Jia, Huang, and Zhang 2008; Guo et al. 2009) , we design 12 Gaussian-MFRAT kernels with the size of 31×31 and θ ranging from 0° to 165°, at intervals of 15° in this paper.

Fig.3 (b) shows that the results of Gaussian-MFRAT are clearer and have less noise. Therefore, we can remove the inefficient response suppression denoising strategy in traditional MFRAT, achieving the simplification and differentiability of the PCEM. By means of PCEM, real palmprint images can be mapped into the PCE domain conveniently.

Transfer from Bézier Palm to PCE Domain

As shown in Fig.4 (a), a generator $G_{B \to P}$ is introduced to transform randomly sampled Bézier curves to PCE images in the first stage. The main structure of the $G_{B \to P}$ is illustrated in Fig.4 (b). To simulate the Gaussian-MFRAT convolution in PCEM, a plug-and-play line feature enhancement block (LFEB) is designed to enforce the network to focus on the line energy features.

The structure of LFEB is also shown in Fig.4 (b). For the input feature $X \in \mathbb{R}^{h \times w \times c}$, the mean value is firstly subtracted on each channel, and then a dilated convolution based on Gaussian-MFRAT kernels are computed. The max-response operation is used to obtain the line energy features. The size of the Gaussian-MFRAT kernel is often large, we thus adopt dilated convolutions to reduce computation and time cost. Finally, the line energy feature is multiplied by a learnable parameter and added to the original feature X to obtain the output. The calculation formula is as follows,

$$f_{LFEB}(X_i) = X_i + s_i f_{MAX}\{w^k_{G-MFRAT} * (X_i - \mu_i)\}_{k=1}^{N_k}, \quad (2)$$

where X_i denotes the feature map on the i-th channel, μ_i denotes the mean of X_i, f_{MAX} represents the max-response operation, $w^k_{G-MFRAT}$ represents the k-th Gaussian-MFRAT kernel, the total number of Gaussian-MFRAT kernels N_k is set

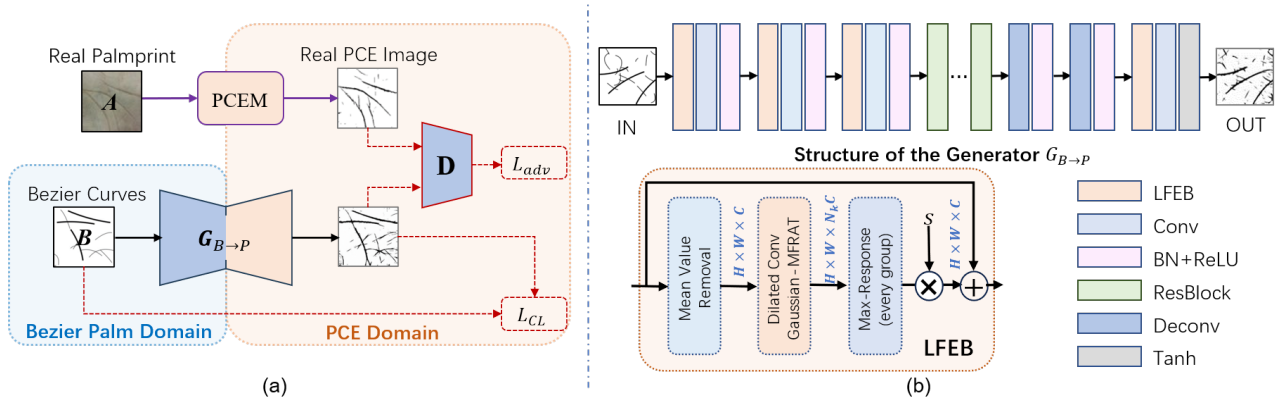

Figure 4: Transfer from Bézier curves to PCE image, (a) framework of the domain transfer model in the first stage, (b) details of the proposed generator $G_{B\to P}$ and the line feature enhancement block (LFEB).

as 12 , and s_i denotes a learnable parameter, which adjusts the degree of feature enhancement for the i-th channel.

The generator $G_{B\to P}$ transforms an input Bézier palm B into a PCE image $G_{B\to P}(B)$. To constrain the structural consistency between B and $G_{B\to P}(B)$, a contrastive loss is used (Park et al. 2020). Specifically, a series of patches are cropped from both B and $G_{B\to P}(B)$. For a query patch q in $G_{B\to P}(B)$, the corresponding patch in B at the same position is positive sample k^+, and other patches in B are negative samples $\{k_i^-\}_{i=1}^N$. Then, the InfoNCE loss (Park et al. 2020) is applied to make the positive samples closer and the negative samples farther apart, as follows:

$$\mathcal{L}_{CL} = \sum_q -\log \frac{\exp(q \cdot k^+/\tau)}{\exp(q \cdot k^+/\tau) + \Sigma_{i=1}^N \exp(q \cdot k_i^-/\tau)}, \quad (3)$$

where τ is a temperature hyper-parameter.

In addition, to optimize the generation quality of PCE images, real PCE images extracted from real palmprints are used as adversarial samples. The commonly used adversarial loss (Goodfellow et al. 2020) is then adopted to drive the generated results closer to real PCE images.

As a result, the total loss in the first stage is computed as:

$$\mathcal{L}_{stage1} = \lambda_1^1 \mathcal{L}_{CL} + \lambda_2^1 \mathcal{L}_{adv}(f_{PCEM}(A), G_{B\to P}(B)), \quad (4)$$

where f_{PCEM} denotes the calculation of PCEM, A represents a random real palmprint image, and λ_1^1, λ_2^1 are two weights.

Generate Realistic Palmprint from PCE Domain

The framework of the second stage is shown in Fig.5, for a real palmprint A, we first map A to the PCE domain using the PCEM. Then the PCE image $f_{PCEM}(A)$ is input to the generator $G_{P\to R}$ to obtain a realistic pseudo palmprint A^* that has consistent palm creases with the input PCE image.

We adopt the same generator backbone as in RPG-palm (Shen et al. 2023), which is a disentanglement-based U-Net structure. The disentanglement-based model can effectively decouple the preservation of crease contents and the generation of styles and details, allowing for diversified outputs by using the latent codes to control styles during inference, which is naturally suitable for few-shot training.

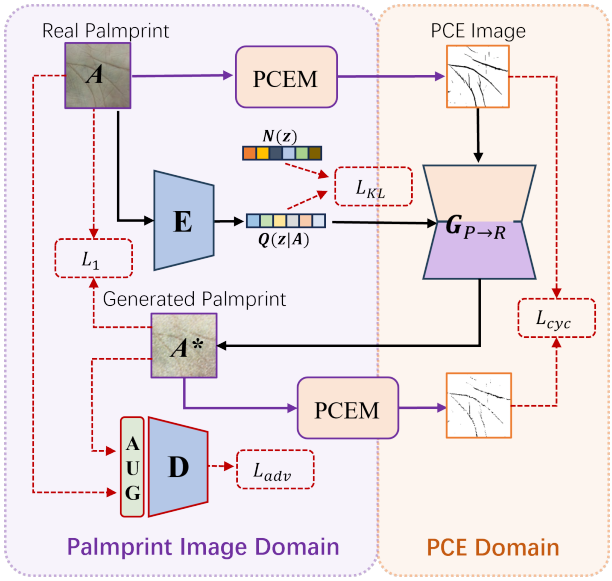

Figure 5: Conditional generation framework from PCE image to the realistic palmprint in the second stage.

The same ResNet-based encoder E (Shen et al. 2023) is used to map the input palmprint A to a latent space $Q(z \mid A) \sim N(\mu_Q, \sigma_Q^2)$ with mean μ_Q and variance σ_Q^2 by reparameterization trick (Zhu et al. 2017b). Then the latent vector is fed into generator $G_{P\to R}$ to control the generated styles. In the training phase, we constrain the divergence of the latent vector to make the distribution of the latent space $N(\mu_Q, \sigma_Q^2)$ approach to $N(0, 1)$. The $\mathcal{L}_{KL}$ loss is calculated as,

$$\mathcal{L}_{KL} = -\frac{1}{2}(1 + \log \sigma_Q^2 - \sigma_Q^2 - \mu_Q^2). \quad (5)$$

By approximating the latent space to the standard normal space, we can easily sample random noise $z \sim N(0, 1)$ as latent control vectors to generate diversified palmprint details in the inference phase.

As mentioned before, it is difficult to train a reliable palmprint recognition network for ID loss (Shen et al. 2023) with

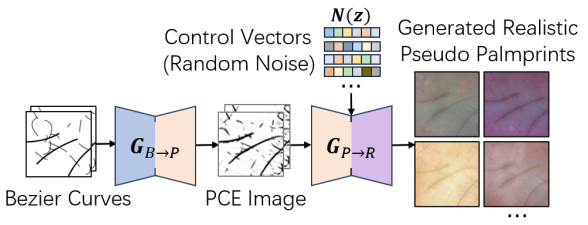

Figure 6: The whole pseudo palmprint generation pipeline in the inference phase.

very few samples. Therefore, we employ a cycled structure with the proposed PCEM to map the generated realistic palmprint A^* back to the PCE domain. By minimizing the $L1$ distance between original $f_{PCEM}(A)$ and generated $f_{PCEM}(A^*)$, the generated palmprint A^* can be constrained to strictly preserve the ID information of the input $f_{PCEM}(A)$. The cycled ID consistency loss is defined as follows,

$$\mathcal{L}_{cyc} = \| f_{PCEM}(A) - f_{PCEM}(G_{P\to R}(f_{PCEM}(A)) \|_1 . \quad (6)$$

In addition, the distortion of the generated palmprint images is constrained by calculating the $\mathcal{L}_1$ loss between palmprint A and the generated palmprint A^*, and a discriminator D based on PatchGAN (Demir and Unal 2018) is applied to enforce the realism of the generated palmprints.

Data augmentation module Note that training with few samples directly may result in overfitting of the discriminator. To mitigate this issue, we adopt a differentiable data augmentation module $\mathcal{T}$ (Karras et al. 2020) before feeding the generated and real images into the discriminator. The generation loss is calculated as,

$$\mathcal{L}_G = \lambda_1^2 \mathcal{L}_1(A, A^*) + \lambda_2^2 \mathcal{L}_{adv}(\mathcal{T}(A), \mathcal{T}(A^*)). \quad (7)$$

The total loss for the second stage is then computed as:

$$\mathcal{L}_{stage2} = \mathcal{L}_G + \lambda_3^2 \mathcal{L}_{KL} + \lambda_4^2 \mathcal{L}_{cyc}, \quad (8)$$

where λ_1^2, λ_2^2, λ_3^2, λ_4^2 denote weights of different loss terms.

Inference Pipeline

The inference process of the proposed method is illustrated in Fig.6. Bézier curves are first generated with randomly sampled control points. Then, the generator $G_{B\to P}$ is used to transfer Bézier curves to the PCE domain. After that, trained generator $G_{P\to R}$ can produce realistic palmprints using PCE images as ID condition. To generate diversified palmprints, random noise vectors are input to $G_{P\to R}$ as the latent control vectors to reproduce various light and skin textures.

Improved BézierPalm

Randomly sampled Bézier curves in BézierPalm (Zhao et al. 2022) are significantly different from the distribution of real palm creases. Therefore, we improve control point sampling mechanism based on prior knowledge from human dermatoglyphics researchs (Park et al. 2010), which divides principle palm lines into 5 types of templates, as shown in Fig.7.

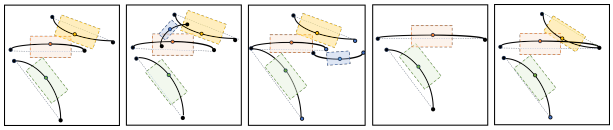

Figure 7: 5 types of improved Bézier curves templates.

Based on the statistical information of human dermatoglyphics researchs, these templates can provide a more accurate range of control points. The improved Bézier curves are closer to the distribution of real creases, thus minimizing the gap between generated and real palmprints.

Improved Recognition Model With LFEB

To further enhance the line features of input palmprints, we tend to apply the proposed LFEB to the recognition model. We incorporate an LFEB before the first convolutional layer of the backbone. This plug-and-play module can enhance the line energy features of the input palmprint images.

Experiment

Experimental Setup

We follow the same experimental datasets and Open-set evaluation protocol as in BézierPalm (Zhao et al. 2022) and RPG-Palm (Shen et al. 2023). The recognition model performance is evaluated in terms of TAR@FAR, where TAR and FAR stand for "True Accept Rate" and "False Accept Rate" respectively. FID (Heusel et al. 2017) metric is used to evaluate the quality of the generated palmprint images.

Datasets We adopt 13 public datasets in our experiments following BézierPalm and RPG-Palm. The public palmprints come from various devices with 3,268 IDs and 59,162 images in total. We follow detect-then-crop protocol (Zhang et al. 2019b) to extract the Region Of Interests (ROIs). Since palm creases are vertically distributed in polyu-ms and polyu-2d+3d, we rotate them 90 degrees clockwise.

Implementation Details of Generation Model Following the BézierPalm and RPG-Palm, we generate 4000 identities with 100 samples for each identity by default. For the first stage, the λ_1^1, λ_2^1, and τ are set to 1.0, 1.0 and 1.0, and the learning rate is 0.0002 in the first 30 epochs and linearly decays to $1e-6$ in the last 30 epochs. For the second stage, we set λ_1^2, λ_2^2, λ_3^2, and λ_4^2 to 1.0, 10.0, 0.01 and 1.0 respectively, and the learning rate is 0.0002 in the first 50 epochs and linearly decays to $1e-8$ in the last 50 epochs according to (Shen et al. 2023). The Adam optimizer parameters are set to $(0.5, 0.99)$ in the training phase. The resolution of all images in the above training is 256×256.

Implementation Details of Recognition Model For a fair comparison, we adopt the same backbones of recognition model as BézierPalm (Zhao et al. 2022), i.e., ResNet50 (He et al. 2016) and MobileFaceNet (Chen et al. 2018) with input resolution of 224×224. The recognition model is firstly pretrained on synthesized data for 25 epochs and then finetuned on real datasets for 50 epochs. The compared baseline model is trained on real datasets for 50 epochs. ArcFace(Deng et al. 2019) with margin $m = 0.5$ and scale fac-

Method	Backbone	train : test = 1 : 1				train : test = 1 : 3			
		TAR@ 1e-3	TAR@ 1e-4	TAR@ 1e-5	TAR@ 1e-6	TAR@ 1e-3	TAR@ 1e-4	TAR@ 1e-5	TAR@ 1e-6
CompCode	N/A	0.4800	0.4292	0.3625	0.2103	0.4501	0.3932	0.3494	0.2648
LLDP	N/A	0.7382	0.6762	0.5222	0.1247	0.7372	0.6785	0.6171	0.2108
BOCV	N/A	0.4930	0.4515	0.3956	0.2103	0.4527	0.3975	0.3527	0.2422
RLOC	N/A	0.6490	0.5884	0.4475	0.1443	0.6482	0.5840	0.5224	0.3366
DOC	N/A	0.4975	0.4409	0.3712	0.1667	0.4886	0.4329	0.3889	0.2007
PalmNet	N/A	0.7174	0.6661	0.5992	0.1069	0.7217	0.6699	0.6155	0.2877
C-LMCL	MB	0.9290	0.8554	0.7732	0.6239	0.8509	0.7554	0.7435	0.5932
ArcFace	MB	0.9292	0.8568	0.7812	0.7049	0.8516	0.7531	0.6608	0.5825
BézierPalm	MB	0.9640	0.9438	0.9102	0.8437	0.9407	0.8861	0.7934	0.7012
RPG-Palm	MB	0.9802	0.9714	0.9486	0.8946	0.9496	0.9267	0.8969	0.8485
Ours	MB	**0.9873**	**0.9806**	**0.9547**	**0.9169**	**0.9674**	**0.9481**	**0.9317**	**0.9079**
C-LMCL	R50	0.9545	0.9027	0.8317	0.7534	0.8601	0.7701	0.6821	0.6254
ArcFace	R50	0.9467	0.8925	0.8252	0.7462	0.8709	0.7884	0.7156	0.6580
BézierPalm	R50	0.9671	0.9521	0.9274	0.8956	0.9424	0.8950	0.8217	0.7649
RPG-Palm	R50	0.9821	0.9732	0.9569	0.9347	0.9533	0.9319	0.9016	0.8698
Ours	R50	**0.9916**	**0.9879**	**0.9827**	**0.9762**	**0.9624**	**0.9626**	**0.9438**	**0.9271**

Table 1: Quantitative results under the open-set protocol where the performances are evaluated in terms of TAR@FAR. 'MB' is MobileFaceNet (Chen et al. 2018) and 'R50' is resnet50 (He et al. 2016).

tor $s = 48$ is used for the pretraining, finetuning, and baseline training supervision. We employ the cosine learning rate schedule with one warmup epoch. Maximum and minimum learning rates for both pretraining and finetuning are set as $1e-2$ and $1e-6$, respectively. All recognition models are trained using a mini-batch SGD with a batch size of 128, and implemented with Pytorch on 4 NVIDIA Tesla V100 GPUs.

Experimental Results

Open-set Palmprint Recognition We first verify the performance of the recognition model under the Open-set protocol, where the training identities and testing identities are completely isolated. We adopt two different training and testing ratios of 1:1 and 1:3 (i.e., train-IDs:test-IDs are $1634 : 1632$ and $818 : 2448$). We compare traditional and deep learning-based palmprint recognition methods (Kong and Zhang 2004; Luo et al. 2016; Guo et al. 2009; Jia, Huang, and Zhang 2008; Fei et al. 2016; Genovese et al. 2019; Zhong and Zhu 2019; Deng et al. 2019; Zhao et al. 2022; Shen et al. 2023), and quantitative results are shown in Tab.1. Our method improves the RPG-Palm with a clear margin and achieves SOTA performance under both 1:1 and 1:3 settings. The improvement of our method under the 1:3 setting is greater than 1:1 setting, showing the effectiveness of our method with fewer real data.

Palmprint Recognition with Limited IDs To validate the performance of our method with limited training data, we test the model with various numbers of training identities under the 1:1 open-set protocol. In our experiments, total of 4000 pseudo IDs are synthesized, and each ID contains 100 pseudo palmprints. The same MobileFaceNet is adopted as a recognition backbone for different methods. As shown in Tab 2, ArcFace, BézierPalm, and RPG-Palm become unusable when training with very few real IDs, but our method is able to maintain the performance. When trained with only 2.5% real IDs (40), our model still outperforms the results

Method	#ID	TAR@FAR= 1e-3	1e-4	1e-5	1e-6
ArcFace	1,600	0.9292	0.8568	0.7812	0.7049
BézierPalm		0.9640	0.9438	0.9102	0.8437
RPG-Palm		0.9802	0.9714	0.9486	0.8946
Ours		**0.9873**	**0.9806**	**0.9547**	**0.9169**
ArcFace	800	0.8934	0.7432	0.7104	0.6437
BézierPalm		0.9534	0.9390	0.9025	0.8164
RPG-Palm		0.9783	0.9687	0.9356	0.8741
Ours		**0.9846**	**0.9718**	**0.9473**	**0.8976**
ArcFace	400	0.8102	0.7050	0.6668	0.3320
BézierPalm		0.9189	0.8497	0.7542	0.6899
RPG-Palm		0.9573	0.9324	0.8836	0.8162
Ours		**0.9679**	**0.9567**	**0.9259**	**0.8572**
ArcFace	160	0.6761	0.5294	0.4783	0.2437
BézierPalm		0.8179	0.6998	0.5826	0.4832
RPG-Palm		0.9356	0.8641	0.8063	0.7246
Ours		**0.9581**	**0.9247**	**0.8861**	**0.8245**
ArcFace	80	0.5384	0.4682	0.3249	0.1173
BézierPalm		0.6547	0.5511	0.4490	0.3743
RPG-Palm		0.8974	0.8092	0.6947	0.5824
Ours		**0.9421**	**0.9009**	**0.8438**	**0.7722**
ArcFace	40	0.4582	0.3908	0.2505	0.0934
BézierPalm		0.6218	0.5145	0.3937	0.3472
RPG-Palm		0.8136	0.6942	0.5894	0.4796
Ours		**0.9351**	**0.8881**	**0.8340**	**0.7694**
ArcFace	16	0.4431	0.3758	0.2372	0.0723
BézierPalm		0.5997	0.4883	0.3542	0.2978
RPG-Palm		0.6958	0.5714	0.4286	0.3271
Ours		**0.9058**	**0.8203**	**0.7469**	**0.6554**

Table 2: Performance under different real training identities. The generation model and the recognition model access the same number of real training identities.

Method	FID↓	train:test=1:1 (with 40 IDs) TAR↑ @1e-3	TAR↑ @1e-4	TAR↑ @1e-5	TAR↑ @1e-6
pix2pixHD	218.1	0.6475	0.5062	0.3708	0.2431
CycleGAN	385.8	0.6982	0.5697	0.4312	0.3162
BicycleGAN	168.8	0.7472	0.6373	0.4978	0.3649
RPG-Palm	188.9	0.8136	0.6942	0.5894	0.4796
Ours	**40.3**	**0.9351**	**0.8881**	**0.8340**	**0.7694**

Table 3: Quantitative recognition results using different generation methods under the open-set protocol.

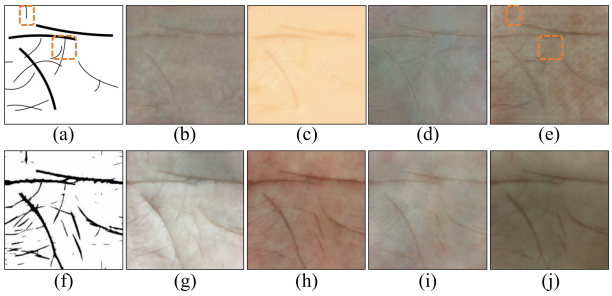

Figure 8: Generated palmprint images of different methods, (a) input BézierPalm, (b) pix2pixHD, (c) CycleGAN, (d) BicycleGAN, (e) RPG-Palm, (f) PCE image, (g)-(j) diversity results of PCE-Palm.

P	A	I	G+L	R+L	TAR@FAR= 1e-4	1e-5	1e-6
✗	✗	✗	✗	✗	0.5503	0.4556	0.4144
✓	✗	✗	✗	✗	0.7022	0.6149	0.5525
✓	✓	✗	✗	✗	0.7449	0.6691	0.6077
✓	✓	✓	✗	✗	0.7757	0.7094	0.6385
✓	✓	✓	✓	✗	0.8238	0.7743	0.7060
✓	✓	✓	✓	✓	**0.8881**	**0.8340**	**0.7694**

Table 4: Ablation of different components in our method. 'P', 'A', 'I', 'G+L', and 'R+L' denote PCEM, DA module, improved BézierPalm, generation model with LFEB, and recognition model with LFEB, respectively.

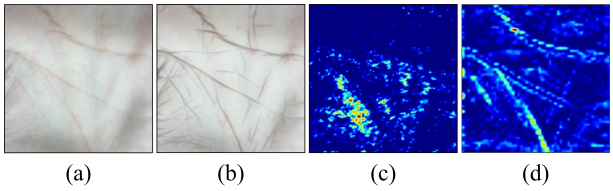

Figure 9: Visualization of input image and features, (a) palmprint image; (b) normalized image after LFEB; (c) feature visualization without LFEB; (d) feature visualization with LFEB.

of ArcFace trained with 100% of real IDs (1600).

Comparison of Generation Methods Four generation methods pix2pixHD (Wang et al. 2018), CycleGAN (Zhu et al. 2017a), BicycleGAN (Zhu et al. 2017b) and RPG-Palm (Shen et al. 2023) are used for comparison, which are all retrained using 40 real IDs with unpaired data following RPG-Palm. The quantitative results are shown in Tab.3. The recognition model pretrained based on our method is much better than other methods. Our method can also achieve an FID score of 40.3, showing a significant improvement. The generated palmprints are illustrated in Fig.8. Under the small amount of training data, the generated samples of RPG-Palm exhibit severe blurring and inconsistent lines, while our method still maintains the overall clarity and ID consistency. Additionally, our method can restore better-detailed information on the palm creases, including the change of thickness. The results of pix2pixHD, CycleGAN, and BicycleGAN exhibit a more severe blur problem than RPG-Palm.

Ablation Study

Components and Design Choices The main components and design choices of our method are PCEM, data augmentation (DA) module for few-samples, improved Bézier curve synthesis, generation model with LFEB, and recognition model with LFEB. We use "P", "A", "I", "G+L" and "R+L" to represent them respectively. For the baseline generation model, we use the two-stage training method and remove the above modules. The ablation experiments are trained using 40 IDs. The test set is fixed under 1:1 open-set protocol. Results of the ablation test are shown in Tab.4. Model with PCEM achieves the greatest performance improvement at 13.81%@FAR=1e-6, which reflects the superiority of generating realistic samples under limited data. By comparing with the baseline generation and recognition model, the LFEB module brings significant and consistent performance gain of 6%@FAR=1e-6 by enhancing palm creases energy features. DA module achieves 5.52%@FAR=1e-6 improvement by effectively expanding the intra-class diversity with few numbers of training samples. In addition, our improvement of Bézier curves also leads to better performance by introducing a more reasonable distribution of palm creases.

Visualizations of LFEB To further verify the effectiveness of the LFEB module, we visualize the features of the layer1 block in MobileFaceNet with and without the enhancement of LFEB in Fig.9. It can be seen in Fig. 9 (c) that the model pays attention to both palm creases and non-line regions without LFEB. By adding the LFEB module, the model can focus on palm creases, as in Fig. 9 (d).

Conclusion

This paper introduces a novel two-stage realistic pseudo palmprint generation method. A new PCE domain is proposed to address the challenge of generating from Bézier to palmprints under limited data. Additionally, a line feature enhancement block is introduced to further improve the generation and recognition models. Furthermore, an improved Bézier curve generation mechanism is proposed, aligning with the distribution of real creases. Experimental results demonstrate that the proposed method significantly outperforms state-of-the-art methods, particularly under few-shot conditions.

Acknowledgements

This work is partly supported by the grants of the National Natural Science Foundation of China, Nos.62076086, 62272142, and 61972129.

References

Almeida, D.; Shmarko, K.; and Lomas, E. 2022. The ethics of facial recognition technologies, surveillance, and accountability in an age of artificial intelligence: a comparative analysis of US, EU, and UK regulatory frameworks. *AI and Ethics*, 2(3): 377–387.

Bahmani, K.; Plesh, R.; Johnson, P.; Schuckers, S.; and Swyka, T. 2021. High fidelity fingerprint generation: Quality, uniqueness, and privacy. In *2021 IEEE International Conference on Image Processing (ICIP)*, 3018–3022. IEEE.

Chen, S.; Liu, Y.; Gao, X.; and Han, Z. 2018. Mobilefacenets: Efficient cnns for accurate real-time face verification on mobile devices. In *Chinese Conference on Biometric Recognition*, 428–438. Springer.

Dai, M.; Hang, H.; and Guo, X. 2022. Adaptive Feature Interpolation for Low-Shot Image Generation. In *European Conference on Computer Vision*, 254–270. Springer.

Demir, U.; and Unal, G. 2018. Patch-based image inpainting with generative adversarial networks. *arXiv preprint arXiv:1803.07422*.

Deng, J.; Guo, J.; Xue, N.; and Zafeiriou, S. 2019. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 4690–4699.

Deng, Y.; Yang, J.; Chen, D.; Wen, F.; and Tong, X. 2020. Disentangled and Controllable Face Image Generation via 3D Imitative-Contrastive Learning. In *IEEE Computer Vision and Pattern Recognition*.

Dian, L.; and Dongmei, S. 2016. Contactless palmprint recognition based on convolutional neural network. In *IEEE ICSP*, 1363–1367. IEEE.

Engelsma, J. J.; Grosz, S. A.; and Jain, A. K. 2022. PrintsGAN: Synthetic fingerprint generator. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Fei, L.; Lu, G.; Jia, W.; Teng, S.; and Zhang, D. 2018. Feature extraction methods for palmprint recognition: A survey and evaluation. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 49(2): 346–363.

Fei, L.; Xu, Y.; Tang, W.; and Zhang, D. 2016. Double-orientation code and nonlinear matching scheme for palmprint recognition. *PR*, 49: 89–101.

Feng, G.; Hu, D.; Zhang, D.; and Zhou, Z. 2006. An alternative formulation of kernel LPP with application to image recognition. *Neurocomputing*, 69(13-15): 1733–1738.

Fu, C.; Wu, X.; Hu, Y.; Huang, H.; and He, R. 2019. Dual variational generation for low shot heterogeneous face recognition. *Advances in neural information processing systems*, 32.

Fu, C.; Wu, X.; Hu, Y.; Huang, H.; and He, R. 2021. Dvg-face: Dual variational generation for heterogeneous face recognition. *IEEE transactions on pattern analysis and machine intelligence*, 44(6): 2938–2952.

Geng, Z.; Cao, C.; and Tulyakov, S. 2019. 3d guided fine-grained face manipulation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9821–9830.

Genovese, A.; Piuri, V.; Plataniotis, K. N.; and Scotti, F. 2019. PalmNet: Gabor-PCA convolutional networks for touchless palmprint recognition. *IEEE TIFS*, 14(12): 3160–3174.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.

Guo, Z.; Zhang, D.; Zhang, L.; and Zuo, W. 2009. Palmprint verification using binary orientation co-occurrence vector. *PRL*, 30(13): 1219–1227.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.

Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.

Hu, D.; Feng, G.; and Zhou, Z. 2007. Two-dimensional locality preserving projections (2DLPP) with its application to palmprint recognition. *PR*, 40(1): 339–342.

Jeong, J.; and Shin, J. 2021. Training gans with stronger augmentations via contrastive discriminator. *arXiv preprint arXiv:2103.09742*.

Jia, W.; Huang, D.-S.; and Zhang, D. 2008. Palmprint verification based on robust line orientation code. *PR*, 41(5): 1504–1513.

Jia, W.; Ren, Q.; Zhao, Y.; Li, S.; Min, H.; and Chen, Y. 2022. EEPNet: An efficient and effective convolutional neural network for palmprint recognition. *Pattern Recognition Letters*, 159: 140–149.

Karras, T.; Aittala, M.; Hellsten, J.; Laine, S.; Lehtinen, J.; and Aila, T. 2020. Training Generative Adversarial Networks with Limited Data. In *Proc. NeurIPS*.

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.

Kong, A.-K.; and Zhang, D. 2004. Competitive coding scheme for palmprint verification. In *ICPR*, volume 1, 520–523. IEEE.

Kong, C.; Kim, J.; Han, D.; and Kwak, N. 2022. Few-shot image generation with mixup-based distance learning. In *European Conference on Computer Vision*, 563–580. Springer.

Li, Z.; Xia, B.; Zhang, J.; Wang, C.; and Li, B. 2022. A comprehensive survey on data-efficient GANs in image generation. *arXiv preprint arXiv:2204.08329*.

Liu, M.-Y.; Huang, X.; Mallya, A.; Karras, T.; Aila, T.; Lehtinen, J.; and Kautz, J. 2019. Few-shot unsupervised image-to-image translation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10551–10560.

Lu, G.; Zhang, D.; and Wang, K. 2003. Palmprint recognition using eigenpalms features. *PRL*, 24(9-10): 1463–1467.

Luo, Y.-T.; Zhao, L.-Y.; Zhang, B.; Jia, W.; Xue, F.; Lu, J.-T.; Zhu, Y.-H.; and Xu, B.-Q. 2016. Local line directional pattern for palmprint recognition. *PR*, 50: 26–44.

Mo, S.; Cho, M.; and Shin, J. 2020. Freeze the discriminator: a simple baseline for fine-tuning gans. *arXiv preprint arXiv:2002.10964*.

Nguyen-Phuoc, T.; Li, C.; Theis, L.; Richardt, C.; and Yang, Y.-L. 2019. Hologan: Unsupervised learning of 3d representations from natural images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7588–7597.

Ojha, U.; Li, Y.; Lu, J.; Efros, A. A.; Lee, Y. J.; Shechtman, E.; and Zhang, R. 2021. Few-shot image generation via cross-domain correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10743–10752.

Park, J. S.; Shin, D. S.; Jung, W.; and Chung, M. S. 2010. Improved analysis of palm creases. *Anatomy & cell biology*, 43(2): 169–177.

Park, T.; Efros, A. A.; Zhang, R.; and Zhu, J.-Y. 2020. Contrastive Learning for Unpaired Image-to-Image Translation. In *European Conference on Computer Vision*.

Qiu, H.; Yu, B.; Gong, D.; Li, Z.; Liu, W.; and Tao, D. 2021. SynFace: Face Recognition with Synthetic Data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10880–10890.

Sang, H.; Yuan, W.; and Zhang, Z. 2009. Research of palmprint recognition based on 2DPCA. In *International Symposium on Neural Networks*, 831–838. Springer.

Shen, L.; Jin, J.; Zhang, R.; Li, H.; Zhang, Y.; Zhang, J.; Ding, S.; Zhao, Y.; and Jia, W. 2023. RPG-Palm: Realistic Pseudo-data Generation for Palmprint Recognition. *ICCV2023 arXiv:2307.14016*.

Shen, L.; Zhang, Y.; Zhao, K.; Zhang, R.; and Shen, W. 2022. Distribution alignment for cross-device palmprint recognition. *Pattern Recognition*, 132: 108942.

Sun, Z.; Tan, T.; Wang, Y.; and Li, S. Z. 2005. Ordinal palmprint represention for personal identification [represention read representation]. In *CVPR*, volume 1, 279–284. IEEE.

Svoboda, J.; Masci, J.; and Bronstein, M. M. 2016. Palmprint recognition via discriminative index learning. In *ICPR*, 4232–4237. IEEE.

Wang, M.; and Ruan, Q. 2006. Palmprint recognition based on two-dimensional methods. In *ICSP*, volume 4. IEEE.

Wang, T.-C.; Liu, M.-Y.; Zhu, J.-Y.; Tao, A.; Kautz, J.; and Catanzaro, B. 2018. High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.

Wyzykowski, A. B. V.; Segundo, M. P.; and de Paula Lemes, R. 2021. Level three synthetic fingerprint generation. In *2020 25th International Conference on Pattern Recognition (ICPR)*, 9250–9257. IEEE.

Yang, C.; Shen, Y.; Xu, Y.; and Zhou, B. 2021. Data-efficient instance generation from instance discrimination. *Advances in Neural Information Processing Systems*, 34: 9378–9390.

Yang, M.; and Wang, Z. 2023. Image Synthesis under Limited Data: A Survey and Taxonomy. *arXiv preprint arXiv:2307.16879*.

Yin, X.; Yu, X.; Sohn, K.; Liu, X.; and Chandraker, M. 2017. Towards Large-Pose Face Frontalization in the Wild. In *In Proceeding of International Conference on Computer Vision*. Venice, Italy.

Zhang, H.; Zhang, Z.; Odena, A.; and Lee, H. 2019a. Consistency regularization for generative adversarial networks. *arXiv preprint arXiv:1910.12027*.

Zhang, Y.; Zhang, L.; Liu, X.; Zhao, S.; Shen, Y.; and Yang, Y. 2019b. Pay by showing your palm: A study of palmprint verification on mobile platforms. In *2019 IEEE International Conference on Multimedia and Expo (ICME)*, 862–867. IEEE.

Zhao, K.; Shen, L.; Zhang, Y.; Zhou, C.; Wang, T.; Zhang, R.; Ding, S.; Jia, W.; and Shen, W. 2022. BézierPalm: A free lunch for palmprint recognition. In *European Conference on Computer Vision*, 19–36. Springer.

Zhao, S.; Liu, Z.; Lin, J.; Zhu, J.-Y.; and Han, S. 2020. Differentiable augmentation for data-efficient gan training. *Advances in neural information processing systems*, 33: 7559–7570.

Zhao, Z.; Singh, S.; Lee, H.; Zhang, Z.; Odena, A.; and Zhang, H. 2021. Improved consistency regularization for gans. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 11033–11041.

Zheng, Q.; Kumar, A.; and Pan, G. 2015. Suspecting less and doing better: New insights on palmprint identification for faster and more accurate matching. *IEEE TIFS*, 11(3): 633–641.

Zhong, D.; and Zhu, J. 2019. Centralized large margin cosine loss for open-set deep Palmprint recognition. *IEEE TCSVT*.

Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017a. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networkss. In *Computer Vision (ICCV), 2017 IEEE International Conference on*.

Zhu, J.-Y.; Zhang, R.; Pathak, D.; Darrell, T.; Efros, A. A.; Wang, O.; and Shechtman, E. 2017b. Toward multimodal image-to-image translation. In *Advances in Neural Information Processing Systems*.