FlowPalm: Optical Flow Driven Non-Rigid Deformation for Geometrically Diverse Palmprint Generation

Supplementary Material

This supplementary material includes:

- detailed algorithms of FlowPalm;
- data selection for generative training;
- additional ablation studies;
- analysis of different numbers of synthetic identities;
- results on additional recognition backbones.

Algorithm 1: Deformation Prior Construction

Input: $\mathcal{D}_{palm}$: real palmprint dataset, RAFT: optical flow estimator, $R(\cdot)$: recognition model, τ_d, τ_c: thresholds
Output: $\mathcal{L}$: deformation library

$\mathcal{L} \leftarrow \emptyset$;
for *each identity* y *in* $\mathcal{D}_{palm}$ **do**
  Sample a palmprint pair (I_s, I_t) from identity y;
  $F \leftarrow \text{RAFT}(I_s, I_t)$;
  Compute flow smoothness $D(F)$;
  if $D(F) < \tau_d$ **then**
    $\hat{I}_t \leftarrow W(I_s, F)$;
    $C(F) \leftarrow \cos\big(R(\hat{I}_t), R(I_t)\big)$;
    if $C(F) > \tau_c$ **then**
      Add F to $\mathcal{L}$;
return $\mathcal{L}$

Algorithm 2: Training with Conditional Dropout

Input: x_0: real palmprint, T_{ext}: texture extractor, p_{drop}: dropout probability
while *not converged* **do**
  Sample a mini-batch $\{x_0\}$ from the training set;
  $C \leftarrow T_{ext}(x_0)$;
  Sample noise $\epsilon \sim \mathcal{N}(0, I)$, timestep $t \sim \mathcal{U}(1, T)$;
  $x_t \leftarrow \sqrt{\bar{\alpha}_t}\, x_0 + \sqrt{1-\bar{\alpha}_t}\, \epsilon$;
  With prob. p_{drop} set $C = \emptyset$;
  $\hat{\epsilon} \leftarrow \epsilon_\theta(x_t, C, t)$;
  $\mathcal{L}_{\text{ddpm}} \leftarrow \|\epsilon - \hat{\epsilon}\|_2^2$;
  Update θ by descending $\nabla_\theta \mathcal{L}_{\text{ddpm}}$;

Algorithm 3: FlowPalm Three-Stage Sampling

Input: C: crease map, $\mathcal{L}$: deformation library, ϵ_θ: trained diffusion model, T: total steps
Output: x_0: generated palmprint

Sample $F \sim \mathcal{L}$, $x_T \sim \mathcal{N}(0, I)$;
Stage I: Deformed Structure;
$C_w \leftarrow W(C, F)$; `// Warp crease with` F
for $t = T$ ***down to*** t^* **do**
  $x_{t-1} \leftarrow \epsilon_\theta(x_t, C_w, t)$;
$x_{\text{clean}} \leftarrow \big(x_{t^*} - \sqrt{1-\bar{\alpha}_{t^*}}\, \epsilon_\theta(x_{t^*}, C_w, t^*)\big) / \sqrt{\bar{\alpha}_{t^*}}$;
Stage II: Texture Generation;
Sample $\xi \sim \mathcal{N}(0, I)$;
$\xi_w \leftarrow T_{warp}(\xi, F)$; `// Warp noise with` F
$x_{t^*} \leftarrow \sqrt{\bar{\alpha}_{t^*}}\, x_{\text{clean}} + \sqrt{1-\bar{\alpha}_{t^*}}\, \xi_w$;
for $t = t^*$ ***down to*** $0.25T$ **do**
  $x_{t-1} \leftarrow \epsilon_\theta(x_t, C_w, t)$;
Stage III: Unconditional Refinement;
Set $C \leftarrow \emptyset$;
for $t = 0.25T$ ***down to*** 0 **do**
  $x_{t-1} \leftarrow \epsilon_\theta(x_t, C, t)$;
return x_0

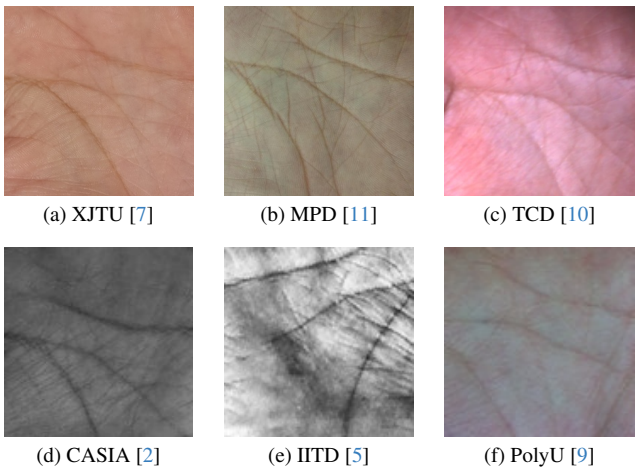

(a) XJTU [7] (b) MPD [11] (c) TCD [10]
(d) CASIA [2] (e) IITD [5] (f) PolyU [9]

Figure 1. Sample images from six benchmark palmprint datasets used in our experiments.

1. Algorithmic Overview

Algorithms 1–3 summarize the three core components of FlowPalm. Algorithm 1 builds a deformation prior by collecting optical flows between real palmprints and filtering them with smoothness and identity-consistency criteria, forming a library of realistic non-rigid deformation. Algorithm 2 trains a diffusion model with conditional dropout, so that it can operate both with texture conditions and without

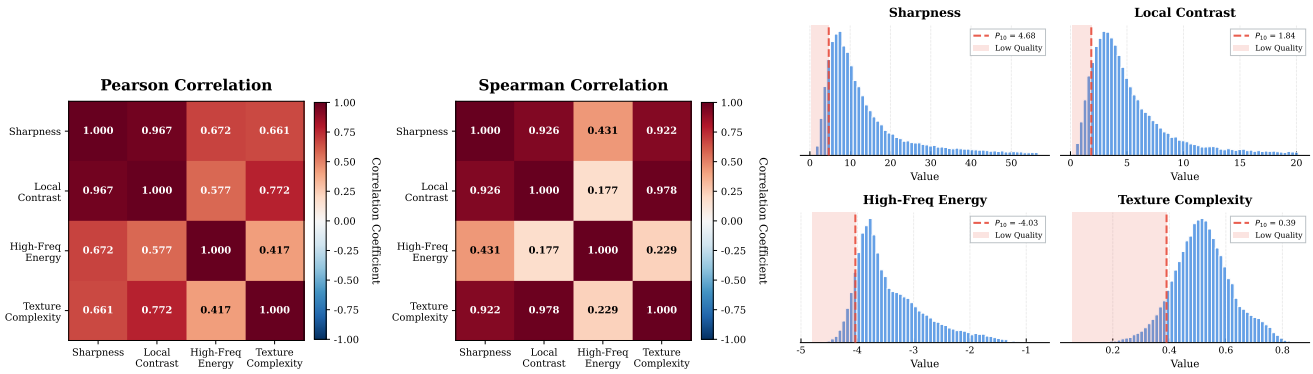

(a) Correlation heatmap

(b) Quality distribution

Figure 2. Visualization of quality feature correlations and their marginal distributions.

Metric	Configs		Score Distributions			Verification Datasets						Avg.
	Clean Step	Stage-II Time	Fréchet Distance↓	Inter-cls Distance↑	Intra-cls Distance↓	XJTU	MPD	TCD	CASIA	IITD	PolyU	
TAR@FAR =1e-3 ↑ (%)	×	–	**0.1408**	0.9558	0.6412	78.09	79.65	97.62	78.87	87.49	98.54	86.71
	✓	$0.1T$	0.1514	0.9503	0.5894	91.89	90.46	99.24	86.80	92.16	99.52	93.34
	✓	$0.3T$	0.1525	0.9545	0.4232	94.54	93.70	99.50	96.00	94.90	**99.69**	96.39
	✓	$0.5T$	0.1503	**0.9559**	0.3766	**95.82**	**94.84**	**99.62**	95.60	**96.11**	99.64	**96.94**
	✓	$0.7T$	0.1507	0.9551	0.3447	94.81	92.39	99.49	95.93	95.33	99.53	96.25
	✓	$0.9T$	0.1507	0.9550	**0.3201**	93.09	90.05	99.55	**97.40**	95.30	99.53	95.82
TAR@FAR =1e-4 ↑ (%)	×	–	**0.1408**	0.9558	0.6412	63.71	64.93	93.25	61.00	78.80	95.34	76.17
	✓	$0.1T$	0.1514	0.9503	0.5894	80.50	79.77	98.14	69.40	87.72	98.65	85.70
	✓	$0.3T$	0.1525	0.9545	0.4232	84.45	85.97	98.97	88.40	91.11	**99.07**	91.33
	✓	$0.5T$	0.1503	**0.9559**	0.3766	**88.51**	**88.04**	**99.21**	83.73	**93.79**	98.57	**91.98**
	✓	$0.7T$	0.1507	0.9551	0.3447	85.06	84.13	98.96	**92.40**	92.29	97.45	91.72
	✓	$0.9T$	0.1507	0.9550	**0.3201**	82.32	78.66	98.95	91.33	91.02	98.90	90.20

Table 1. Additional ablation on the clean step and Stage-II start time. The clean step indicates whether a clean denoising step is performed as the end step in Stage I.

conditions (for unconditional refinement). Algorithm 3 then defines a three-stage sampling procedure: Stage I shapes a deformed but identity-consistent ridge structure, Stage II injects deformation-consistent textures via warped noise, and Stage III performs unconditional refinement to remove artifacts and better match real palmprint statistics.

2. Training Data for Generative Model

As illustrated in Figure 1, the six benchmark datasets exhibit diverse imaging conditions and quality levels. To obtain high quality inputs for training the generative model, we apply a quality filtering strategy. For each ROI, we compute four scalar features: sharpness, local contrast, high-frequency energy, and texture complexity. Their correlations and marginal distributions are shown in Figure 2, from which we derive a 10th-percentile threshold for each feature on real data. An image is regarded as high quality if at least three of the four features exceed their respective thresholds (i.e., at most one feature falls below the 10% quantile); otherwise it is assigned to the low-quality set. The high-quality subset is then used as the main training data for our generative model.

3. Additional Ablation Studies

Table 1 examines how the clean step in Stage I and the starting time of Stage II affect both distribution metrics and downstream verification performance. We compare a baseline without the clean step against several configurations that enable it while varying the Stage-II start time. The results show that introducing the clean step significantly improves intra-class similarity. Starting Stage II earlier further increases intra-class similarity but simultaneously reduces intra-class diversity. Setting the Stage-II start time to around $0.5T$ achieves the best overall performance.

Metric	Training Data	IDs	Verification Datasets						Avg.
			XJTU	MPD	TCD	CASIA	IITD	PolyU	
TAR@FAR =1e-3 ↑ (%)	Real Data	–	88.05	94.71	98.19	68.93	95.43	99.30	90.77
	FlowPalm Data	500	84.28	83.89	98.79	85.93	93.07	99.22	90.86
		1,000	91.01	90.76	99.46	94.80	93.56	99.58	94.86
		2,000	95.82	94.84	99.62	95.60	**96.11**	**99.64**	96.94
		4,000	96.41	95.94	99.81	97.20	95.56	99.62	97.42
		8,000	**98.06**	**98.00**	**99.92**	**97.73**	95.39	99.58	**98.11**
TAR@FAR =1e-4 ↑ (%)	Real Data	–	76.92	84.13	95.98	38.53	91.47	98.54	80.93
	FlowPalm Data	500	66.94	74.65	96.71	72.00	88.89	98.27	82.91
		1,000	75.56	81.01	98.70	88.07	89.19	98.41	88.49
		2,000	88.51	88.04	99.21	83.73	**93.79**	98.57	91.98
		4,000	90.67	89.50	99.50	91.27	93.07	**99.11**	93.85
		8,000	**93.95**	**93.35**	**99.69**	**94.20**	93.04	98.76	**95.50**

Table 2. Verification performance on six benchmark datasets under different numbers of training identities. All models are trained with RandAugment-based augmentation [1].

Backbone	Metric	Method	Venue	Verification Datasets						Avg.
				XJTU	MPD	TCD	CASIA	IITD	PolyU	
ResNet18	TAR@FAR =1e-3 ↑ (%)	RPG-Palm [8]	ICCV'23	39.76	57.65	87.02	45.87	91.08	95.14	69.42
		PCE-Palm [3]	AAAI'24	38.27	55.85	78.64	46.13	90.13	73.94	63.83
		Diff-Palm [4]	CVPR'25	64.41	79.14	95.58	75.60	94.58	94.58	83.98
		PFIG-Palm [12]	TIP'25	66.88	77.75	95.95	67.73	93.83	97.65	83.30
		FlowPalm (Ours)	–	**90.87**	**90.02**	**98.93**	**92.13**	**96.44**	**99.08**	**94.58**
	TAR@FAR =1e-4 ↑ (%)	RPG-Palm [8]	ICCV'23	23.63	43.57	75.42	28.53	85.59	87.77	57.42
		PCE-Palm [3]	AAAI'24	23.91	42.04	66.90	28.00	84.61	59.85	50.89
		Diff-Palm [4]	CVPR'25	45.29	64.08	89.98	59.47	90.36	87.90	72.85
		PFIG-Palm [12]	TIP'25	32.78	62.71	90.32	50.60	89.48	93.68	69.93
		FlowPalm (Ours)	–	**81.11**	**80.58**	**97.54**	**86.13**	**94.28**	**98.15**	**89.63**
ResNet50	TAR@FAR =1e-3 ↑ (%)	RPG-Palm [8]	ICCV'23	43.22	56.98	88.22	61.80	91.11	94.16	72.58
		PCE-Palm [3]	AAAI'24	45.94	61.81	81.48	44.67	89.87	80.56	67.39
		Diff-Palm [4]	CVPR'25	61.64	81.56	97.00	72.67	95.30	96.84	84.17
		PFIG-Palm [12]	TIP'25	67.20	78.97	96.90	62.07	93.50	98.05	82.78
		FlowPalm (Ours)	–	**92.70**	**92.32**	**99.43**	**93.87**	**96.80**	**99.48**	**95.77**
	TAR@FAR =1e-4 ↑ (%)	RPG-Palm [8]	ICCV'23	25.72	41.92	78.20	44.87	85.30	86.53	60.42
		PCE-Palm [3]	AAAI'24	28.09	50.17	70.52	29.67	84.02	70.10	55.43
		Diff-Palm [4]	CVPR'25	39.46	68.46	93.17	57.80	92.42	92.53	73.97
		PFIG-Palm [12]	TIP'25	41.98	64.23	91.36	44.80	88.76	94.40	70.92
		FlowPalm (Ours)	–	**84.01**	**83.40**	**98.51**	**87.40**	**94.90**	**98.64**	**91.14**

Table 3. Quantitative comparison of different generation methods across additional recognition backbones. All models are trained using RandAugment-based augmentation [1].

4. Effect of the Number of Synthetic Identities

Table 2 analyzes how the number of training identities influences verification performance when using FlowPalm-generated data. We compare a model trained on real data with models trained solely on FlowPalm data, where the number of synthetic identities is gradually increased from 500 to 8,000. Notably, using FlowPalm data with only 500 synthetic identities already yields an average performance that surpasses the real-data baseline. As the number of synthetic identities increases, the TAR at both FAR=10^{-3} and FAR=10^{-4} consistently improves, demonstrating that FlowPalm scales effectively to large identity sets.

5. Additional Recognition Backbones

To verify that the benefits of FlowPalm are not restricted to a specific recognition architecture, Table 3 reports additional results on two widely used backbones, ResNet18 and ResNet50. For both TAR@FAR=10^{-3} and TAR@FAR=10^{-4}, FlowPalm consistently achieves the highest average performance across all six datasets on both backbones.

6. Limitations

Although FlowPalm achieves strong performance, it still has an inherent limitation. The deformation priors in our framework are obtained from optical flows estimated between real palmprint pairs. While these flows do not contain identity-specific information and are only used to model general non-rigid patterns, this reliance on real data may still restrict scalability.

In future work, we plan to investigate synthetic deformation priors, for example by learning a generative model capable of producing plausible flow fields or by computing deformations on 3D simulated hand models [6]. Such directions would further improve privacy and the general applicability of FlowPalm.

References

[1] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, pages 702–703, 2020. 3

[2] Ying Hao, Zhenan Sun, Tieniu Tan, and Chao Ren. Multispectral palm image fusion for accurate contact-free palmprint recognition. In *Proc. IEEE Int. Conf. Image Process. (ICIP)*, pages 281–284, 2008. 1

[3] Jianlong Jin, Lei Shen, Ruixin Zhang, Chenglong Zhao, Ge Jin, Jingyun Zhang, Shouhong Ding, Yang Zhao, and Wei Jia. Pce-palm: Palm crease energy based two-stage realistic pseudo-palmprint generation. In *Proc. AAAI Conf. Artif. Intell. (AAAI)*, pages 2616–2624, 2024. 3

[4] Jianlong Jin, Chenglong Zhao, Ruixin Zhang, Sheng Shang, Jianqing Xu, Jingyun Zhang, ShaoMing Wang, Yang Zhao, Shouhong Ding, Wei Jia, et al. Diff-palm: Realistic palmprint generation with polynomial creases and intra-class variation controllable diffusion models. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 26367–26376, 2025. 3

[5] Ajay Kumar and Sumit Shekhar. Personal identification using multibiometrics rank-level fusion. *IEEE Trans. Syst. Man Cybern.*, 41:743–752, 2010. 1

[6] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, 36(6):245:1–245:17, 2017. 4

[7] Huikai Shao, Dexing Zhong, and Xuefeng Du. Deep distillation hashing for unconstrained palmprint recognition. *IEEE Trans. Instrum. Meas.*, 70:1–13, 2021. 1

[8] Lei Shen, Jianlong Jin, Ruixin Zhang, Huaen Li, Kai Zhao, Yingyi Zhang, Jingyun Zhang, Shouhong Ding, Yang Zhao, and Wei Jia. Rpg-palm: Realistic pseudo-data generation for palmprint recognition. In *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, pages 19605–19616, 2023. 3

[9] David Zhang, Zhenhua Guo, Guangming Lu, Lei Zhang, and Wangmeng Zuo. An online system of multispectral palmprint verification. *IEEE Trans. Instrum. Meas.*, 59:480–490, 2009. 1

[10] Lin Zhang, Lida Li, Anqi Yang, Ying Shen, and Meng Yang. Towards contactless palmprint recognition: A novel device, a new benchmark, and a collaborative representation based identification approach. *Pattern Recognit.*, 69:199–212, 2017. 1

[11] Yingyi Zhang, Lin Zhang, Ruixin Zhang, Shaoxin Li, Jilin Li, and Feiyue Huang. Towards palmprint verification on smartphones. *arXiv preprint arXiv:2003.13266*, 2020. 1

[12] Yuchen Zou, Huikai Shao, Chengcheng Liu, Siyu Zhu, Zongqing Hou, and Dexing Zhong. Pfig-palm: Controllable palmprint generation via pixel and feature identity guidance. *IEEE Trans. Image Process.*, 34:6603–6615, 2025. 3