GENERATING SYNTHETIC FACE RECOGNITION DATASETS USING BROWNIAN IDENTITY DIFFUSION AND A FOUNDATION MODEL

Hatef Otroshi Shahreza[1,2] *and Sébastien Marcel*[1,3]

[1]Idiap Research Institute, Switzerland
[2]École Polytechnique Fédérale de Lausanne (EPFL), Switzerland
[3]Université de Lausanne (UNIL), Switzerland

{hatef.otroshi,sebastien.marcel}@idiap.ch

ABSTRACT

Training face recognition models requires a large amount of identity-labeled face images, which are often collected by crawling the web, and therefore have ethical and privacy concerns. Recently, generating synthetic face datasets and training face recognition models using synthetic datasets has emerged to be a viable solution. This paper presents BIF-Face, a new framework to generate synthetic face recognition datasets. We use the Brownian identity diffusion to generate synthetic identities, and then build synthetic face recognition datasets by generating different samples per each identity using a foundation model. In our experiments, we use the generated face datasets to train face recognition models and evaluate them on several real benchmarking dataset. Our experimental results show that face recognition models trained with BIF-Face achieve competitive performance with face recognition models trained on state-of-the-art synthetic face recognition datasets.

Index Terms— Face Recognition, Synthetic Dataset, Brownian Identity Diffusion, Foundation Model

1. INTRODUCTION

The advancements in state-of-the-art face recognition models are derived apart from angular loss functions [1, 2] and availability of large-scale datasets [3, 4]. However, existing large-scale face recognition datasets, such as MS-Celeb [3], WebFace260M [4], etc., are collected without individuals' consent, raising ethical, legal, and privacy concerns. With the growing debates on the privacy of using web-crawled datasets, several of such large-scale face recognition datasets, such as MS-Celeb [3], VGGFace2 [5], etc., have been retracted, to prevent potential legal issues. These circumstances have created doubts regarding the continuity of research and development for face recognition, since the availability of data, which is such a fundamental resource, has become uncertain.

Fig. 1: Sample images from our BIF-Face synthetic dataset.

With the recent advancements in generative models, generating synthetic face recognition datasets has recently emerged as a promising alternative solution to address the legal and ethical concerns with real face recognition datasets [6–11]. The generated synthetic dataset requires to have different synthetic identities with different samples per each identity. Bae *et al.* [12] used a computer-graphic pipeline to render different identities and synthesize different images for each identity. In contrast, most papers in the literature used Generative Adversarial Networks (GANs) or Diffusion Models (DMs) to generate synthetic datasets. Qiu *et al.* [13] used DiscoFaceGAN [14] to generate different synthetic identities and used identity mixup by exploring the latent space of DiscoFaceGAN to increase intra-class variation. Boutros *et al.* [15] trained an identity-conditioned StyleGAN2 [16] on the CASIA-WebFace [17] dataset, and then generated a synthetic face recognition dataset using the trained model. Kolf *et al.* [18] also trained an identity-conditioned StyleGAN2 [16] in a three-player GAN framework to integrate the identity information into the generation process. Colbois

This work was funded by the Hasler foundation through the Responsible Face Recognition (SAFER) project, the H2020 TReSPAsS-ETN Marie Skłodowska-Curie early training network (agreement 860813), and the Swiss Center for Biometrics Research & Testing at Idiap Research Institute.

979-8-3503-2411-2/25/$31.00 ©2025 IEEE

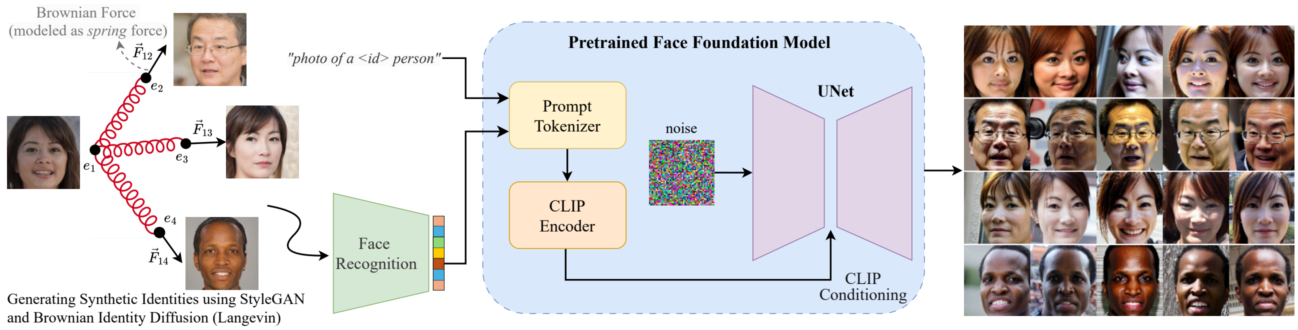

Fig. 2: Block diagram of our proposed framework to generate synthetic datasets. We first generate synthetic identities using StyleGAN and then generate different samples for each identity using a pretrained face foundation model.

et al. [19] used a pretrained StyleGAN2 [16] and trained a support vector machine (SVM) to find directions for different variations in the face image. Then, they used StyleGAN2 to generate different identities and generated different samples for each identity by exploring the intermediate latent space of StyleGAN2. Geissbuhler *et al.* [20] sampled different identities in the latent space of pretrained StyleGAN2 [16]. Inspired by the physical motion of soft particles subjected to stochastic Brownian forces, they modeled identities as Brownian particles and used the dynamical equation (called Langavien) to find different identities. Then, they solved a similar equation in the latent space of StyleGAN2 to generate different samples per identity (called Dispersion) and further explored the intermediate latent space (called DisCo).

In contrast to most work which used GAN-based generator models, a few works used diffusion models. Kim *et al.* [21] trained a dual condition (style and identity conditions) face generator diffusion model on CASIA-WebFace. Then they used their trained model to generate different identities and different styles for each identity. Boutros *et al.* [22] used an identity-conditioned latent diffusion model to generate synthetic face images. They generated different samples by an unconditional diffusion model and then generated different samples using their conditional latent diffusion model (called IDiff-Face Two-Stage method). Alternatively, they uniformly sampled identity from the hypersphere and generated different samples for each identity using their identity-conditioned latent diffusion model (called IDiff-Face Uniform method). Melzi *et al.* [23] proposed a hybrid dataset generation method. They used StyleGAN to generate face images with different identities and then used DreamBooth [24] as a diffusion-based generator, to generate different samples for each identity.

In this paper, we propose a new hybrid framework to generate synthetic datasets using StyleGAN and a face foundation model based on diffusion models. We use the Brownian identity diffusion [20] to generate different synthetic identities using StyleGAN, by solving the dynamical equation for randomly generated identities as soft particles. Then, we consider the generated synthetic identities as reference images and generate different samples for each identity using a face foundation model [25] based on Stable Diffusion [26]. Fig 1 illustrates sample images of our synthetic dataset. In our experiments, we generate several synthetic face recognition datasets and train face recognition models with our generated datasets. We evaluate the trained face recognition models on several real benchmarking datasets. Our experimental results show that the face recognition models trained on our synthetic datasets achieve competitive results with models trained on state-of-the-art synthetic datasets.

The remainder of this paper is organized as follows. In Section 2, we describe our framework to generate synthetic datasets using Brownian identity diffusion and a face foundation model. In Section 3 we provide our experimental results and evaluate our synthetic datasets. Finally, we conclude the paper in Section 4.

2. PROPOSED FRAMEWORK

In the first step in our hybrid framework for generating synthetic datasets, we use StyleGAN [16] to generate different synthetic identities using the Brownian identity diffusion [20]. Then, we generate different sample images for each identity using a face foundation model [25]. In the following sections, we describe each step separately. Fig. 2 illustrates the general block diagram of our framework.

2.1. Synthesizing Identities

To generate synthetic identities, we use StyleGAN [16] and explore its latent space using the Langevin algorithm [20]. The Langevin algorithm starts with random sampling from StyleGAN latent space and considers different identities as soft particles in a solvent that have stochastic Brownian forces. By modeling the Brownian forces between particles with spring force and iteratively solving the dynamical equation for the Brownian particle, the Langevin algorithm can find different synthetic identities. The generated images by

Table 1: Comparison of recognition performance of face recognition models trained with different synthetic and real datasets. The performance reported for each dataset is in terms of accuracy and best values for each dataset type are emboldened.

Dataset Type	Dataset name	N_{id}	$N_{samples}$	LFW	CPLFW	CALFW	CFP	AgeDB
Real images	MS-Celeb-1M [3]	85'000	5'800'000	**99.82**	92.83	**96.07**	96.10	**97.82**
	WebFace-4M [4]	206'000	4'000'000	99.78	**94.17**	95.98	**97.14**	97.78
	CASIA-WebFace [17]	10'572	490'623	99.42	90.02	93.43	94.97	94.32
Computer Graphics	DigiFace-1M [12]	109'999	1'219'995	90.68	72.55	73.75	79.43	68.43
Diffusion-based	DCFace-0.5M [21]	10'000	500'000	98.35	83.12	91.70	88.43	89.50
	DCFace-1.2M [21]	60'000	1'200'000	**98.90**	**84.97**	**92.80**	**89.04**	**91.52**
	IDiff-Face (Uniform) [22]	10'049	502'450	98.18	80.87	90.82	82.96	85.50
	IDiff-Face (Two-Stage) [22]	10'050	502'500	98.00	77.77	88.55	82.57	82.35
GAN-based	Synface [13]	10'000	999'994	86.57	65.10	70.08	66.79	59.13
	SFace [15]	10'572	1'885'877	93.65	74.90	80.97	75.36	70.32
	Syn-Multi-PIE [19]	10'000	180'000	78.72	60.22	61.83	60.84	54.05
	IDnet [18]	10'577	1'057'200	84.48	68.12	71.42	68.93	62.63
	Langevin-Dispersion [20]	10'000	650'000	94.38	65.75	86.03	65.51	77.30
	Langevin-DisCo [20]	10'000	650'000	97.07	76.73	89.05	79.56	83.38
	Langevin-DisCo [20]	30'000	1'950'000	**98.97**	**81.52**	**93.95**	**83.77**	**93.32**
Hybrid Generation	GANDiffFace [23]	10'080	543'893	94.35	76.15	79.90	78.99	69.82
	BIF-Face [ours]	10'000	640'000	98.10	80.97	88.28	84.43	84.18
	BIF-Face [ours]	30'000	1'920'000	**98.43**	**83.58**	**90.35**	**88.76**	**88.45**

StyleGAN using the Langevin algorithm are considered as reference images for different identities in our framework.

2.2. Synthesizing Different Samples for Each Identity

We consider the images generated by StyleGAN as reference images and use a face foundation model to generate different samples for each identity. Foundation models have gained significant attention in recent years [27]. We use Arc2Face [25] as the foundation model, in which CLIP [28] and Stable Diffusion [26] models were fine-tuned on the WebFace260M dataset [4]. The CLIP encoder is fine-tuned to transform the text encoder into a face recognition embedding space, so that with the prompt of *"photo of a $< id >$ person"* the model can generate a face image given face recognition embedding for the $< id >$ token. Therefore, the resulting model takes as input face embeddings of a predefined face recognition model along with a random noise and can generate different face images for each embedding. Hence, we first extract face embeddings from reference images generated using Langevin algorithm by StyleGAN, and then generate different samples using the foundation model by changing the noise input. As shown in Fig 1, the generated images for each identity have different variations, including light condition, pose, etc.

3. EXPERIMENTS

3.1. Experimental Setup

To evaluate the performance of generated synthetic datasets, we use the generated datasets to train a face recognition model with the same configuration. We use the iResNet50 backbone and train it with AdaFace loss function [2] for 30 epochs using the Stochastic Gradient Descent (SGD) optimizer with the initial learning rate 0.1 and weight decay 5×10^{-4}. We benchmark the performance of the trained models on different real datasets, including Labeled Faces in the Wild (LFW) [29], Cross-age LFW (CA-LFW) [30], Cross-Pose LFW (CP-LFW) [31], Celebrities in Frontal-Profile in the Wild (CFP-FP) [32], AgeDB-30 [33], and IARPA Janus Benchmark-C (IJB-C) [34] datasets. For consistency with prior works, we report accuracy calculated using 10-fold cross-validation on LFW, CA-LFW, CP-LFW, CFP-FP, and AgeDB datasets. However, for the IJB-C dataset we consider True Accept Rate (TAR) at different False Accept Rate (FAR) values, and report Receiver Operating Characteristic (ROC) plots. The source code of our experiments and generated datasets are publicly available[1]..

3.2. Comparison

To compare the generated dataset with our framework to previous synthetic face recognition datasets, we train differ-

[1]https://gitlab.idiap.ch/biometric/code.paper.mlsp2025_biface

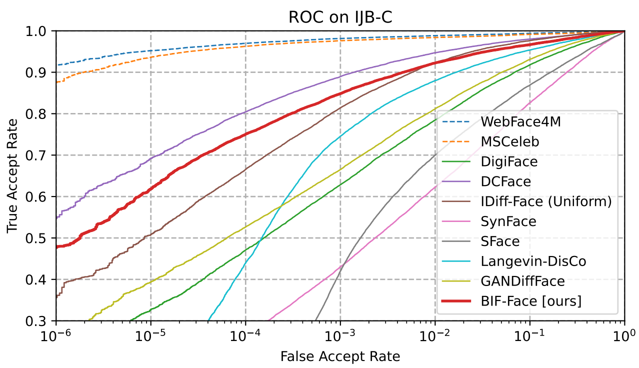

Fig. 3: Recognition performance on IJB-C dataset.

Table 2: Ablation study on the effect of number of samples per identity

$N_{samples}$	LFW	CPLFW	CALFW	CFP	AgeDB
32	98.07	80.12	88.3	82.87	83.38
42	98.22	80.40	87.73	**84.79**	84.17
50	**98.40**	80.38	88.03	83.86	**84.40**
64	98.10	**80.97**	**88.28**	84.43	84.18
96	98.15	80.30	88.08	84.43	83.30
128	98.07	79.35	87.70	83.74	83.45

ent face recognition models with the same configuration as described in Section 3.1 and evaluate the recognition performance on benchmarking datasets. Table 1 compares the recognition performance of different models trained with synthetic datasets as well as three different real datasets. As the results in this table show our method achieves comparable performance with state-of-the-art synthetic datasets in the literature. Fig. 3 also compares the ROC of face recognition models on the IJB-C dataset. To simplify the ROC plots and have a fair comparison, we plot ROC curves for versions of each synthetic dataset with 10K identities. The plots in this figure show that our proposed method achieves superior performance than GAN-based and hybrid datasets and is competitive with state-of-the-art diffusion-based datasets. In particular, the face recognition model trained with our synthetic dataset outperforms the models trained with IDiff-Face and achieves comparable performance with the one trained with DCFace. However, identities generated by Brownian identity diffusion using Langevin algorithm are shown to have less leakage of identity compared to previous methods, such as DCFace [20, 35]. We can also observe that the model trained with our hybrid dataset also outperforms the one trained with GANDiffFace, which is the only hybrid synthetic dataset in the literature, on all benchmarks. Comparing the recognition performance with real datasets in Table 1 and Fig. 3, we can also observe that there is still a gap in the performance of face recognition models trained with synthetic and real datasets.

Table 3: Ablation study on the effect of number of identities

N_{id}	LFW	CPLFW	CALFW	CFP	AgeDB
10K	98.10	80.97	88.28	84.43	84.18
20K	**98.50**	81.93	89.70	86.63	87.05
30K	98.43	**83.58**	**90.35**	88.76	**88.45**
40K	98.03	82.38	89.93	88.87	87.93
50K	98.28	83.40	90.05	**89.09**	87.42

Table 4: Ablation study on the effect of backbones

Backbone	LFW	CPLFW	CALFW	CFP	AgeDB
IResNet18	97.8	77.5	87.95	80.23	82.93
IResNet34	98.17	80.93	**88.73**	83.2	82.92
IResNet50	98.3	80.58	87.63	84.1	83.88
IResNet101	**98.33**	**82.25**	88.63	**85.51**	**85.23**

3.3. Ablation Study

We perform several ablation studies to evaluate the effect of different hyperparameters on the performance of the generated datasets. In our first ablation study, we consider 10k synthetic identities and generate different numbers of images per identity. Table 2 compares the performance of the face recognition model trained on synthetic datasets with different numbers of samples per identity. As the results in this table show increasing the number of samples per identity improves the recognition performance and then saturates or slightly degrades the performance.

As our second ablation study, we generate different datasets with different numbers of identities using the Langevin method [20], and then generate 64 samples for each set of identities. Table 3 reports the recognition performance of the face recognition model trained on synthetic datasets with different numbers of identities. Similar to Table 2, we can observe that increasing the number of identities leads to a higher recognition performance, but the recognition performance is saturated and does not further improve.

As another ablation study, we consider different backbones and evaluate the recognition performance of face recognition models trained on our synthetic dataset with 10K identities and 64 samples per identity. Table 4 compares the recognition performance of IResNet backbones with different numbers of layers. As the results show, an IResNet model with a larger number of layers achieves a better performance.

4. CONCLUSION

In this paper, we proposed a new framework, called BIF-Face, to generate hybrid synthetic face recognition datasets (based on GAN and diffusion models). We used StyleGAN to generate different synthetic identities using Brownian identity diffusion and then generated different samples for each identity using a foundation model. In our experiments, we used the generated datasets to train different face recognition models and evaluated the trained models on several real benchmarking datasets. We also compared the performance of trained face recognition models using our synthetic dataset and previous synthetic datasets in the literature. The experimental results show that our synthetic datasets achieve competitive performance compared to state-of-the-art synthetic face recognition datasets.

5. REFERENCES

[1] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4690–4699.

[2] Minchul Kim, Anil K Jain, and Xiaoming Liu, "Adaface: Quality adaptive margin for face recognition," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 18750–18759.

[3] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao, "Ms-celeb-1m: A dataset and benchmark for large-scale face recognition," in *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part III 14*. Springer, 2016, pp. 87–102.

[4] Zheng Zhu, Guan Huang, Jiankang Deng, Yun Ye, Junjie Huang, Xinze Chen, Jiagang Zhu, Tian Yang, Jiwen Lu, Dalong Du, et al., "Webface260m: A benchmark unveiling the power of million-scale deep face recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 10492–10502.

[5] Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman, "Vggface2: A dataset for recognising faces across pose and age," in *2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018)*. IEEE, 2018, pp. 67–74.

[6] Hatef Otroshi Shahreza et al., "Sdfr: Synthetic data for face recognition competition," in *2024 IEEE 18th International Conference on Automatic Face and Gesture Recognition (FG)*. IEEE, 2024, pp. 1–9.

[7] Pietro Melzi et al., "Frcsyn-ongoing: Benchmarking and comprehensive evaluation of real and synthetic data to improve face recognition systems," *Information Fusion*, vol. 107, pp. 102322, 2024.

[8] Hatef Otroshi Shahreza and Sébastien Marcel, "Hyperface: Generating synthetic face recognition datasets by exploring face embedding hypersphere," in *The Thirteenth International Conference on Learning Representations*, 2025.

[9] Ivan DeAndres-Tame, Ruben Tolosana, Pietro Melzi, Ruben Vera-Rodriguez, Minchul Kim, Christian Rathgeb, Xiaoming Liu, Luis F Gomez, Aythami Morales, Julian Fierrez, et al., "Second frcsyn-ongoing: Winning solutions and post-challenge analysis to improve face recognition with synthetic data," *Information Fusion*, vol. 120, pp. 103099, 2025.

[10] Parsa Rahimi, Behrooz Razeghi, and Sebastien Marcel, "Synthetic to authentic: Transferring realism to 3d face renderings for boosting face recognition," in *European Conference on Computer Vision*. Springer, 2024, pp. 109–126.

[11] Hatef Otroshi Shahreza, Anjith George, and Sébastien Marcel, "Synthdistill: Face recognition with knowledge distillation from synthetic data," in *2023 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 2023, pp. 1–10.

[12] Gwangbin Bae, Martin de La Gorce, Tadas Baltrušaitis, Charlie Hewitt, Dong Chen, Julien Valentin, Roberto Cipolla, and Jingjing Shen, "Digiface-1m: 1 million digital face images for face recognition," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 3526–3535.

[13] Haibo Qiu, Baosheng Yu, Dihong Gong, Zhifeng Li, Wei Liu, and Dacheng Tao, "Synface: Face recognition with synthetic data," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 10880–10890.

[14] Yu Deng, Jiaolong Yang, Dong Chen, Fang Wen, and Xin Tong, "Disentangled and controllable face image generation via 3d imitative-contrastive learning," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 5154–5163.

[15] Fadi Boutros, Marco Huber, Patrick Siebke, Tim Rieber, and Naser Damer, "Sface: Privacy-friendly and accurate face recognition using synthetic data," in *2022 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 2022, pp. 1–11.

[16] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila, "Analyzing and improving the image quality of stylegan," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 8110–8119.

[17] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li, “Learning face representation from scratch,” *arXiv preprint arXiv:1411.7923*, 2014.

[18] Jan Niklas Kolf, Tim Rieber, Jurek Elliesen, Fadi Boutros, Arjan Kuijper, and Naser Damer, “Identity-driven three-player generative adversarial network for synthetic-based face recognition,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 806–816.

[19] Laurent Colbois, Tiago de Freitas Pereira, and Sébastien Marcel, “On the use of automatically generated synthetic image datasets for benchmarking face recognition,” in *2021 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 2021, pp. 1–8.

[20] David Geissbühler, Hatef Otroshi Shahreza, and Sébastien Marcel, “Synthetic face datasets generation via latent space exploration from brownian identity diffusion,” in *Forty-second International Conference on Machine Learning*, 2025.

[21] Minchul Kim, Feng Liu, Anil Jain, and Xiaoming Liu, “Dcface: Synthetic face generation with dual condition diffusion model,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 12715–12725.

[22] Fadi Boutros, Jonas Henry Grebe, Arjan Kuijper, and Naser Damer, “Idiff-face: Synthetic-based face recognition through fizzy identity-conditioned diffusion model,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 19650–19661.

[23] Pietro Melzi, Christian Rathgeb, Ruben Tolosana, Ruben Vera-Rodriguez, Dominik Lawatsch, Florian Domin, and Maxim Schaubert, “Gandiffface: Controllable generation of synthetic datasets for face recognition with realistic variations,” *arXiv preprint arXiv:2305.19962*, 2023.

[24] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman, “Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 22500–22510.

[25] Foivos Paraperas Papantoniou, Alexandros Lattas, Stylianos Moschoglou, Jiankang Deng, Bernhard Kainz, and Stefanos Zafeiriou, “Arc2face: A foundation model of human faces,” in *European Conference on Computer Vision*, 2024.

[26] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer, “High-resolution image synthesis with latent diffusion models,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 10684–10695.

[27] Hatef Otroshi Shahreza and Sébastien Marcel, “Foundation models and biometrics: A survey and outlook,” *Authorea Preprints*, 2025.

[28] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al., “Learning transferable visual models from natural language supervision,” in *International conference on machine learning*. PMLR, 2021, pp. 8748–8763.

[29] Gary B Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller, “Labeled faces in the wild: A database forstudying face recognition in unconstrained environments,” in *Workshop on faces in'Real-Life'Images: detection, alignment, and recognition*, 2008.

[30] Tianyue Zheng, Weihong Deng, and Jiani Hu, “Cross-age lfw: A database for studying cross-age face recognition in unconstrained environments,” *arXiv preprint arXiv:1708.08197*, 2017.

[31] Tianyue Zheng and Weihong Deng, “Cross-pose lfw: A database for studying cross-pose face recognition in unconstrained environments,” *Beijing University of Posts and Telecommunications, Tech. Rep*, vol. 5, no. 7, 2018.

[32] Soumyadip Sengupta, Jun-Cheng Chen, Carlos Castillo, Vishal M Patel, Rama Chellappa, and David W Jacobs, “Frontal to profile face verification in the wild,” in *2016 IEEE winter conference on applications of computer vision (WACV)*. IEEE, 2016, pp. 1–9.

[33] Stylianos Moschoglou, Athanasios Papaioannou, Christos Sagonas, Jiankang Deng, Irene Kotsia, and Stefanos Zafeiriou, “Agedb: the first manually collected, in-the-wild age database,” in *proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2017, pp. 51–59.

[34] Brianna Maze, Jocelyn Adams, James A Duncan, Nathan Kalka, Tim Miller, Charles Otto, Anil K Jain, W Tyler Niggel, Janet Anderson, Jordan Cheney, et al., “Iarpa janus benchmark-c: Face dataset and protocol,” in *2018 international conference on biometrics (ICB)*. IEEE, 2018, pp. 158–165.

[35] Hatef Otroshi Shahreza and Sébastien Marcel, “Unveiling synthetic faces: How synthetic datasets can expose real identities,” *arXiv preprint arXiv:2410.24015*, 2024.