https://doi.org/10.1051/itmconf/20268403028

Fingerprint Image Generation Using Deep Generative Models: From GANs to Diffusion Models

Boyu Zheng[1*]

[1]North University of China, Jiancaoping District, Taiyuan, Shanxi 030051, China

Abstract. This paper reviews recent progress in fingerprint image generation using deep generative models, with a focus on Generative Adversarial Network (GAN)-based and diffusion-based approaches. Fingerprint data are essential for biometric recognition, but collecting large and diverse datasets is difficult, especially for latent fingerprints. Early methods based on physical or statistical modeling could not produce realistic textures or sufficient diversity. With the development of deep learning, GAN-based models such as FingerGAN, PrintsGAN, and lightweight GANs have significantly improved the realism of generated fingerprints by learning data distributions directly. These methods introduce techniques such as structural constraints, multi-stage generation, and improved loss functions to enhance image quality and stability. However, GAN-based models still suffer from problems such as training instability, mode collapse, and limited control over identity consistency. To address these issues, diffusion models have recently been introduced into fingerprint generation. By gradually denoising random noise, diffusion models can generate high-quality and diverse fingerprint images with better stability. Advanced diffusion frameworks further enable controllable generation, allowing users to adjust fingerprint attributes such as style, quality, and sensor type while preserving identity information. Overall, diffusion-based methods show strong potential to become the next generation of fingerprint synthesis techniques.

1 Introduction

Fingerprints are one of the most common biometric traits used for identity verification. They have unique ridge patterns that do not change over time, which makes them reliable for many applications, such as unlocking phones, access control, banking security and forensic investigation [1-3]. A fingerprint image contains clear ridge–valley structures and small local details called minutiae, and these features help fingerprint recognition systems work accurately. Because of this, having enough high-quality fingerprint data is very important.

However, collecting large fingerprint datasets is not easy. Many public datasets contain only a small number of people and only a few impressions for each finger. Latent fingerprints,

* Corresponding author's email: 2205040735@st.nuc.edu.cn

which are prints left unintentionally on surfaces, are even harder to collect. They are often blurred, incomplete, or noisy, and their collection may involve privacy concerns. Several commonly used datasets have also been restricted in recent years. As a result, the amount of fingerprint data available for research and model training is limited, especially for tasks that require high diversity and large sample sizes.

Before deep learning became popular, researchers tried to generate synthetic fingerprints using physical or statistical models [4, 5]. These methods simulate ridge flow, global shape, and minutiae patterns. Although useful, they often produce images that look artificial, because it is difficult to hand-design all the fine textures seen in real fingerprints. Their generated ridge lines may also look too smooth or too regular. These problems limit the use of such methods in real-world fingerprint applications.

The development of deep learning brought new ideas for fingerprint generation. The introduction of Generative Adversarial Networks (GANs) made it possible to create more realistic fingerprint images by learning the data distribution directly. Several fingerprint GAN models have been proposed. For example, FingerGAN generates enhanced latent fingerprints using a constrained GAN framework [6]. Models such as Finger-GAN also focus on generating ridge textures with better structure and continuity [7]. Lightweight GAN models allow faster training and large-scale generation with reduced artifacts [8]. In addition, PrintsGAN can generate multiple impressions of the same synthetic finger, which helps capture natural variations inside the same identity [9]. These GAN-based methods improve realism and diversity compared with earlier approaches.

Even with these improvements, GANs still have some problems. Their training process is unstable and may suffer from mode collapse, where the generator produces limited variations. GANs can also be sensitive to training settings, and their results may not always be consistent, especially for high-resolution fingerprint images.

In recent years, diffusion models have become a strong alternative. Diffusion models generate images by gradually removing noise step by step. This process is more stable and often produces images with better diversity and detail. Studies on latent fingerprint synthesis show that diffusion models can produce clearer ridge patterns, better global structure, and more natural variations than GANs [10]. More advanced designs, such as the GenPrint controllable diffusion model, allow users to control fingerprint class, sensor type, image quality, and acquisition style while keeping identity consistent across multiple impressions [11]. This makes diffusion models useful for both research and training fingerprint recognition systems.

Because fingerprint generation is becoming an important topic, and both GANs and diffusion models are widely used today, a clear summary of existing work is needed. The goal of this paper is to give an overview of fingerprint image generation methods based on these two types of deep generative models. The paper first introduces the background of fingerprint recognition and the problem of data scarcity, then explains basic ideas of GANs and diffusion models, then reviews representative GAN-based and diffusion-based fingerprint generation studies, and finally discusses current challenges, future directions, and concludes the survey.

2 Preliminaries of generative adversarial network and diffusion model

Deep generative models aim to learn the data distribution of real images and then generate new samples from this learned distribution. Two widely used families of generative models are the GAN and the diffusion model. This section gives a simple overview of their basic ideas.

2.1 Generative adversarial network

A Generative Adversarial Network consists of two neural networks: a generator G and a discriminator D [12]. The generator takes a random noise vector z and tries to produce an image that looks real. The discriminator tries to tell real images from fake images generated by G.

Both networks are trained together in a game-like process. The generator improves by trying to fool the discriminator, and the discriminator improves by detecting fake images more accurately. The original GAN objective is a minimax problem:

$$\min_G \max_D \ \mathbb{E}_{x\sim p_{data}(x)}[\log D(x)] + \mathbb{E}_{z\sim p_z(z)}[\log(1 - D(G(z)))] \tag{1}$$

Here, $p_{data}$is the distribution of real images, and $p_z$is the noise distribution.
The goal of G is to generate images $G(z)$ that match the real image distribution, so that $D(G(z)) \approx 1$.

In practice, many variants of GANs are used to make training easier and more stable, such as DCGAN, WGAN, and conditional GAN. For fingerprint applications, GANs are used to generate ridge patterns, enhance latent fingerprints, or synthesize multiple impressions for one identity. Although GANs can produce sharp images, they may suffer from unstable training and mode collapse, where the generator only produces limited patterns.

2.2 Diffusion model

Diffusion models generate images through a two-step process: 1) A forward process that gradually adds noise to an image. 2) A reverse process that learns to remove noise step by step to recover clean images [13].

In the forward process, a real image $x_0$is transformed into a noisy image x_t through T steps:

$$q(x_t \mid x_0) = \mathcal{N}(x_t; \sqrt{\alpha_t}x_0,\ (1-\alpha_t)I) \tag{2}$$

where α_t controls how much noise is added. After many steps, x_T becomes almost pure Gaussian noise.

The reverse process is learned using a neural network ϵ_θ, which predicts the noise added at each step:

$$p_\theta(x_{t-1} \mid x_t) = \mathcal{N}(x_{t-1};\ \mu_\theta(x_t, t),\ \sigma_t^2 I) \tag{3}$$

By repeatedly denoising from x_T back to x_0, the model generates a clean synthetic image. Diffusion models are stable to train and can produce highly detailed and diverse images. Recent works show that diffusion models perform especially well on fingerprint generation tasks, including latent fingerprint synthesis and controllable fingerprint generation.

3 The application of generative adversarial network in fingerprint generation

3.1 FingerGAN

FingerGAN is a GAN-based method designed for latent fingerprint enhancement. The key idea is to treat enhancement as a constrained fingerprint generation problem [6]. Instead of directly improving a noisy latent print, the model generates a clean version while following correct fingerprint structure. To achieve this, FingerGAN introduces two important types of prior information: a fingerprint skeleton map and a FOMFE-based orientation field. These two maps describe ridge connectivity and ridge direction, which help guide the generator to produce more realistic ridge flow. The generator in FingerGAN is a U-shaped network (U-Net). The encoder extracts features from the latent fingerprint, and the decoder reconstructs a clean and complete ridge map. Skip connections help preserve fine details. The discriminator receives not only the real or generated image, but also the skeleton map and orientation field. This allows the discriminator to check both realism and structural correctness, pushing the generator to follow real fingerprint geometry.

3.2 PrintsGAN

PrintsGAN is a novel framework for generating realistic synthetic fingerprint images, aiming to solve the lack of large-scale public fingerprint datasets [9]. The key innovation lies in how the method models both inter-class diversity (different fingers) and intra-class variation (multiple impressions of the same finger), which previous approaches failed to achieve.

Unlike earlier GAN-based methods that directly generate fingerprints from noise, PrintsGAN divides the generation process into three structured stages. First, it generates a Master-Print, which represents the core identity of a finger. Second, it applies a learned nonlinear warping module to simulate natural variations caused by finger pressure, rotation, and placement during acquisition. Third, a texture rendering module adds realistic ridge details using a style-based GAN, producing visually convincing fingerprints. This design allows the model to generate multiple realistic impressions for the same identity, which is essential for training recognition systems. Another key contribution is that PrintsGAN integrates domain knowledge from traditional fingerprint modeling with modern GAN architectures. This hybrid design allows better control over fingerprint structure and realism compared with purely data-driven methods. Extensive experiments shown in Fig. 1 show that fingerprints generated by PrintsGAN are visually close to real ones and achieve high realism scores in expert evaluations.

Fig. 1. Generated fingerprint images based on PrintsGAN [9].

3.3 A lightweight GAN

A lightweight GAN-based framework was proposed for generating high-quality synthetic fingerprint images [8]. The main goal is to address common problems in fingerprint generation, such as unstable training, mode collapse, and low image diversity. To achieve this, the authors design a new GAN architecture that combines several effective strategies in a simple and efficient way.

First, the model introduces spectral normalization in key layers of the generator and discriminator. This helps control the Lipschitz constant of the network and stabilizes training, preventing sudden gradient explosions. Second, the authors design an average residual connection instead of standard residual addition. This structure helps avoid vanishing or exploding gradients and improves convergence during training. Third, the paper proposes a novel loss doping strategy, where a small adaptive perturbation is added to the generator loss when training stagnates. This prevents the generator from getting stuck and encourages continuous learning. The proposed network is lightweight and can generate realistic fingerprint images up to 256×256 resolution. Compared with existing methods such as DCGAN, WGAN, and BEGAN, it produces clearer ridge patterns and higher diversity, measured using the MS-SSIM metric. Experimental results shown in Fig. 2 show that the method achieves better visual quality and diversity while using fewer parameters.

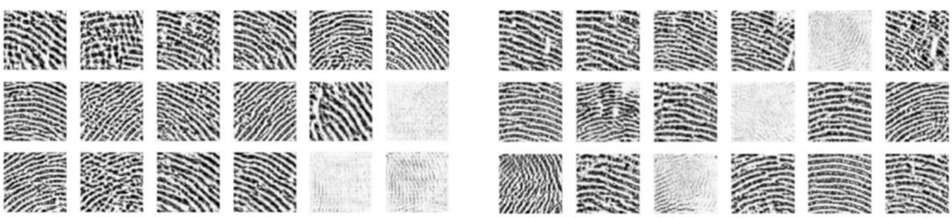

Fig. 2. Generated fingerprint images based on the lightweight GAN [8].

4 The application of diffusion model in fingerprint generation

4.1 Vanilla diffusion probabilistic models

This paper studies how to generate realistic latent fingerprint images using diffusion probabilistic models (DDPMs) [10]. Latent fingerprints are often incomplete, noisy, and hard to collect, which makes it difficult to train reliable fingerprint recognition systems. To solve this problem, the authors explore whether diffusion models can generate high-quality and diverse synthetic fingerprint images.

The paper first reviews traditional fingerprint generation methods, including physical models, statistical models, and GAN-based approaches. Although GANs can generate realistic images, they often suffer from training instability, mode collapse, and limited diversity. To overcome these problems, the authors apply a diffusion probabilistic model, which generates images through a gradual denoising process. This approach allows the model to learn the data distribution more stably and accurately.

The proposed method is evaluated using a large dataset of 10,000 latent fingerprint images. Experimental results show that the diffusion model produces more realistic and diverse fingerprints than GAN-based methods such as DCGAN and StyleGAN. Both quantitative metrics (FID scores) and human visual evaluations confirm its superior performance. The generated fingerprints also preserve fine ridge details and structural patterns more effectively.

4.2 Controllable diffusion model

This paper proposes a new fingerprint generation method based on diffusion probabilistic models, aiming to improve both the quality and controllability of synthetic fingerprint images [11]. Traditional fingerprint generation methods, including GAN-based models, often suffer from unstable training and limited control over identity and appearance. To address these issues, the authors introduce a diffusion-based framework that can generate realistic fingerprints while allowing fine control over different fingerprint attributes.

The key innovation of this work is the use of a controllable diffusion model that separates identity information from appearance variations. The model introduces a structure-aware representation that preserves the core ridge pattern of a fingerprint while allowing flexible changes in style, such as sensor type, noise level, and image quality. By combining a diffusion backbone with an identity-preserving guidance mechanism, the method can generate multiple realistic fingerprint samples from the same identity without losing structural consistency. This design enables both high diversity and strong identity consistency, which are difficult to achieve simultaneously in GAN-based approaches.

In addition, the method supports multimodal conditions, allowing the generation process to be guided by different types of inputs, such as fingerprint type or capture conditions. Experimental results show that the proposed approach produces higher-quality and more diverse fingerprint images compared to existing GAN-based methods. Overall, this work demonstrates that diffusion models offer a powerful and flexible solution for controllable fingerprint synthesis and provide strong potential for future biometric research and data augmentation tasks.

5 Discussion

5.1 Challenges

Despite the rapid progress in fingerprint image generation using deep generative models, several challenges still remain. One major difficulty lies in data quality and diversity. Although diffusion models and GANs can generate realistic fingerprints, the training data itself is often limited, imbalanced, or noisy. Many datasets contain fingerprints collected under similar conditions, which restricts the model's ability to generalize to different sensors, environments, or capture styles. This limitation is especially serious for latent fingerprints, which often contain blur, distortion, and partial ridge information.

Another challenge is identity consistency. While diffusion models can generate visually realistic fingerprints, maintaining consistent identity across multiple generated samples remains difficult. Some models produce high-quality images but fail to preserve fine ridge structures that define individual identity, which limits their use in biometric recognition systems. Balancing diversity and identity preservation is still an open problem.

Computational cost is also a major concern. Diffusion models require long training times and high computational resources due to their iterative denoising process. This makes them less suitable for real-time or large-scale deployment compared to GAN-based models. In addition, model training and inference often require powerful GPUs, which may limit accessibility.

Another challenge lies in evaluation. Current evaluation metrics, such as FID or visual inspection, do not fully reflect fingerprint usability in real biometric systems. There is still a lack of standardized metrics that can accurately measure identity preservation, ridge clarity, and forensic usefulness.

Finally, ethical and privacy concerns must be considered. While synthetic data helps reduce privacy risks, highly realistic fingerprints could also be misused if not properly controlled. Ensuring responsible use and secure deployment remains an important challenge.

5.2 Future prospects

Future research on fingerprint generation is expected to move toward more controllable, efficient, and reliable models. One promising direction is the development of hybrid frameworks that combine diffusion models with lightweight architectures or GAN components, aiming to reduce computational cost while maintaining high image quality. Improving sampling efficiency and accelerating inference will be critical for real-world applications.

Another important direction is fine-grained controllability. Future models may allow explicit control over fingerprint attributes such as ridge density, orientation, pressure patterns, and sensor-specific noise. This would make generated data more suitable for targeted tasks like algorithm testing, data augmentation, and robustness evaluation.

The integration of domain knowledge, such as ridge flow rules and fingerprint anatomy, is also expected to improve realism and interpretability. Combining data-driven learning with domain constraints can help models generate fingerprints that are not only visually convincing but also biologically meaningful.

Additionally, future research may focus on multi-modal and cross-domain learning, where fingerprint generation is combined with other biometric modalities or sensor conditions. This could improve generalization and enable broader forensic applications.

Finally, as diffusion models continue to evolve, their application to fingerprint synthesis is likely to expand toward controllable, efficient, and trustworthy generation frameworks, supporting both research and practical biometric systems.

6 Conclusion

In summary, fingerprint image generation has evolved from traditional modeling methods to advanced deep generative models. GAN-based approaches greatly improved realism but still face challenges in stability and controllability. Diffusion models provide a promising alternative by offering higher image quality, stronger diversity, and better control over generation. Although challenges such as computational cost and identity preservation remain, diffusion-based fingerprint generation represents a powerful and flexible direction for future biometric research and practical applications.

References

1. D. Maltoni, D. Maio, A. K. Jain, and S. Prabhakar, Handbook of Fingerprint Recognition (Springer, London, 2009).
2. M. M. Ali, V. H. Mahale, P. Yannawar, and A. T. Gaikwad, Overview of fingerprint recognition system, in Proceedings of the 2016 International Conference on Electrical, Electronics, and Optimization Techniques (ICEEOT) (IEEE, 2016), pp. 1334–1338.
3. A. K. Jain and A. Kumar, Biometric recognition: An overview, in Second Generation Biometrics: The Ethical, Legal and Social Context (Springer, 2012), pp. 49–79.
4. R. Cappelli, D. Maltoni, D. Maio, A. K. Jain, and S. Prabhakar, Synthetic fingerprint generation, in Handbook of Fingerprint Recognition (Springer, 2003), pp. 203–232.

5. R. Cappelli, D. Maio, and D. Maltoni, Synthetic fingerprint-database generation, in Proceedings of the 2002 International Conference on Pattern Recognition (IEEE, 2002), Vol. 3, pp. 744–747.
6. Y. Zhu, X. Yin, and J. Hu, FingerGAN: A constrained fingerprint generation scheme for latent fingerprint enhancement, IEEE Transactions on Pattern Analysis and Machine Intelligence 45, 8358–8371 (2023).
7. S. Minaee and A. Abdolrashidi, Finger-GAN: Generating realistic fingerprint images using connectivity imposed GAN, arXiv preprint arXiv:1812.10482 (2018).
8. M. A. Fahim and H. Y. Jung, A lightweight GAN network for large-scale fingerprint generation, IEEE Access 8, 92918–92928 (2020).
9. J. J. Engelsma, S. Grosz, and A. K. Jain, PrintsGAN: Synthetic fingerprint generator, IEEE Transactions on Pattern Analysis and Machine Intelligence 45, 6111–6124 (2023).
10. J. Wang, Z. Zhang, and C. Han, Exploring latent fingerprint synthesis with diffusion probabilistic models, in Proceedings of the Future Technologies Conference (Springer, Cham, 2024), pp. 78–89.
11. S. A. Grosz and A. K. Jain, Universal fingerprint generation: Controllable diffusion model with multimodal conditions, IEEE Transactions on Pattern Analysis and Machine Intelligence (2024).
12. A. Creswell, T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta, and A. A. Bharath, Generative adversarial networks: An overview, IEEE Signal Processing Magazine 35, 53–65 (2018).
13. F. A. Croitoru, V. Hondru, R. T. Ionescu, and M. Shah, Diffusion models in vision: A survey, IEEE Transactions on Pattern Analysis and Machine Intelligence 45, 10850–10869 (2023).