Adversarially Robust Fingerprint Authentication via PGD-Aware Training and Diffusion-Based Purification

Alan Wilson[1]

[1] Intact Financial Corporation, Toronto, Ontario M5G 0A1, Canada

Correspondence: Alan Wilson, Intact Financial Corporation, Toronto, Ontario M5G 0A1, Canada.

Received: May 8, 2026 Accepted: June 4, 2026 Published: June 8,2026
doi:10.65343/aiis.v2i1.97 URL: https://doi.org/10.65343/aiis.v2i1.97

Abstract

Fingerprint recognition is widely used for biometric authentication in digital financial systems, but its vulnerability to adversarial perturbations is a growing security concern. Projected Gradient Descent (PGD) attacks are particularly damaging—by iteratively refining perturbations within a constrained noise budget, they can cause fingerprint recognition models to fail in ways that simpler one-step attacks cannot. In this paper, we propose a two-stage defense that pairs PGD-aware adversarial training with a diffusion model-based purification module (DiffPure). The first stage exposes the model to PGD-perturbed samples during training, building intrinsic robustness. The second stage runs a score-based diffusion model at inference time to clean adversarial noise from incoming fingerprint images before recognition. We use EfficientNet-B3 as the backbone and cosine similarity for identity matching. On the SOCOFing dataset, the combined framework reduces EER under PGD-7 attack from 0.42 to 0.19—a 55% relative improvement—while keeping clean-input EER at 0.23. Neither defense alone comes close to this. The results suggest that adversarial training and diffusion purification are genuinely complementary, and that pairing them is worth the added complexity.

Keywords: fingerprint recognition, adversarial robustness; PGD attack, diffusion model, biometric authentication

1. Introduction

Fingerprint recognition has become one of the most widely deployed biometric modalities in digital financial systems. Banks, mobile payment platforms, and access control systems have all moved toward fingerprint-based authentication as a way to provide fast, contactless identity verification without requiring users to remember passwords or carry tokens (Li, Li, & Chen, 2021; Jain, Ross, & Prabhakar, 2004). Compared to other biometric modalities—iris texture, facial geometry, voice—fingerprints offer a practical combination of high discriminability, lifelong stability, and low acquisition cost (Wang, Zhang, & Zhao, 2022; Maltoni, Maio, Jain, & Prabhakar, 2009). Regulatory pressure has reinforced adoption: the EU's Payment Services Directive 2 (PSD2), for instance, mandates strong customer authentication, and fingerprint recognition has become a common technical answer to that requirement.

The security picture is less reassuring. Fingerprint recognition systems are vulnerable to adversarial attacks—deliberately crafted perturbations that are imperceptible to the human eye but cause recognition models to fail (Goodfellow, Shlens, & Szegedy, 2015; Akhtar, & Mian, 2018). In the fingerprint context, these perturbations can take two forms: impersonation attacks, where a fraudster's modified fingerprint is accepted as a registered user's, and evasion attacks, where a legitimate user is incorrectly rejected (Engelsma, & Jain, 2022; Bhatt, Bharadwaj, Singh, & Vatsa, 2015). Both are serious threats in financial authentication, and neither requires the attacker to physically replicate a fingerprint.

Among adversarial attack methods, Projected Gradient Descent (PGD) (Madry, Makelov, Schmidt, Tsipras, & Vladu, 2018) is particularly difficult to defend against. Unlike the Fast Gradient Sign Method (FGSM) (Goodfellow, Shlens, & Szegedy, 2015), which applies a single gradient step, PGD iterates over multiple steps within a constrained perturbation budget, producing much stronger adversarial examples. Our earlier work showed that FGSM attacks significantly degrade face recognition performance and that convolutional autoencoders can partially recover it (Ma, & Wilson, 2023). But autoencoders are not designed with PGD in mind, and the stronger iterative structure of PGD calls for more robust countermeasures.

Defenses against adversarial attacks in biometric systems generally fall into two camps. Model-level defenses modify the training process to build in robustness—adversarial training (Madry, Makelov, Schmidt, Tsipras, & Vladu, 2018; Shafahi, Najibi, Ghiasi, Xu, Dickerson, Studer, ... Goldstein, 2019) being the most established, where the model is trained on adversarial examples generated on the fly. Input-level defenses instead preprocess incoming samples to strip out perturbations before the model ever sees them; spatial smoothing, feature squeezing, and

autoencoder reconstruction (Ma, & Wilson, 2023) all work this way. Each approach has real limitations. Adversarial training alone leaves residual vulnerability to strong iterative attacks and adds substantial training cost. Input-level defenses can hurt clean performance when the perturbation type at test time differs from what the defense was designed for.

Score-based diffusion models (Ho, Jain, & Abbeel, 2020; Song, Sohl-Dickstein, Kingma, Kumar, Ermon, & Poole, 2021) have opened an interesting new option for input-level purification. The core idea behind DiffPure (Nie, Guo, Huang, Xiao, Vahdat, & Anandkumar, 2022) is to run a few steps of the diffusion forward process on an adversarial input—adding just enough noise to disrupt adversarial structure—then run the reverse process to recover a clean image on the natural data manifold. Because the purification works by projecting back onto what the model has learned about clean data, it does not require any prior knowledge of the attack strategy. DiffPure has performed well against PGD and AutoAttack in image classification benchmarks, but its application to fingerprint recognition for financial authentication has not been studied.

This paper proposes a two-stage defense combining PGD-aware adversarial training with DiffPure for fingerprint recognition in digital financial systems. Stage one builds robustness into the EfficientNet-B3 backbone through PGD adversarial training. Stage two deploys a diffusion purification module at inference time to preprocess incoming fingerprint images before recognition. The two stages are complementary by design: purification reduces the adversarial noise burden on the model, and adversarial training handles whatever residual perturbations survive purification. Experiments on the SOCOFing dataset show the combined framework outperforms either defense alone by a meaningful margin.

The rest of the paper is organized as follows. Section 2 reviews related work on adversarial attacks, diffusion-based purification, and fingerprint recognition architectures. Section 3 describes the proposed framework. Section 4 presents experimental results. Section 5 concludes and outlines future directions.

2. Literature Review

2.1 Adversarial Attacks on Biometric Systems

The adversarial vulnerability of deep networks was first documented by Szegedy et al. (Szegedy, Zaremba, Sutskever, Bruna, Erhan, Goodfellow, & Fergus, 2014), who found that small, structured input modifications could reliably cause misclassification with high confidence. Goodfellow et al. (Goodfellow, Shlens, & Szegedy, 2015) formalized this with FGSM—a single gradient step in the direction of increasing loss, fast to compute but relatively weak. Madry et al. (Madry, Makelov, Schmidt, Tsipras, & Vladu, 2018) strengthened this into PGD by iterating the gradient step multiple times while projecting back onto the L-infinity perturbation ball at each iteration, producing adversarial examples that are substantially harder to defend against.

Adversarial attacks have been studied across fingerprint (Ghiani, Yambay, Mura, Tocco, Marcialis, Roli, & Schuckers, 2013; (Engelsma, Cao, & Jain, 2021)), face (Ma, & Wilson, 2023; Ma, & Wilson, 2023), iris (Fang, Czajka, & Bowyer, 2021), and gait modalities (Delgado-Escano, Castro, Cozar, Marin-Jimenez, Guil, & de la Torre, 2020). In fingerprints, ridge-valley pattern encoders are sensitive to pixel-level noise, and studies have shown that PGD-perturbed images can fool both classical minutiae matchers and deep learning-based systems (Ghiani, Yambay, Mura, Tocco, Marcialis, Roli, & Schuckers, 2013). In face recognition, our earlier work showed that FGSM attacks substantially raise EER and that a convolutional autoencoder reconstruction approach partially offsets this—though it also made clear that stronger iterative attacks would require more capable defenses (Ma, & Wilson, 2023). A domain adaptation framework for face recognition under challenging lighting conditions (Ma, & Wilson, 2023) drew on related ideas about distributional robustness, suggesting a broader design principle that carries into the fingerprint adversarial setting.

Several defense strategies have been explored beyond adversarial training. Defensive distillation (Papernot, McDaniel, Wu, Jha, & Swami, 2016) trains a compressed model on soft probability outputs to smooth decision boundaries. Input transformation approaches (Xie, Wang, Zhang, Ren, & Yuille, 2018) apply randomization or filtering at inference time to disrupt adversarial structure. Certified defenses (Cohen, Rosenfeld, & Kolter, 2019) offer provable robustness guarantees within bounded perturbation radii. Each involves trade-offs: distillation and certified methods can sacrifice clean accuracy, while input transformations vary in effectiveness depending on attack type. Adversarial training (Madry, Makelov, Schmidt, Tsipras, & Vladu, 2018) remains empirically one of the strongest defenses when the training and test attacks are matched, but it does not fully close the door on iterative attacks that push beyond the training budget. Hybrid approaches that pair adversarial training with purification have shown promise (Nie, Guo, Huang, Xiao, Vahdat, & Anandkumar, 2022) and motivate the framework proposed here.

2.2 Diffusion Models for Adversarial Purification

Diffusion models learn a forward process that gradually corrupts data with Gaussian noise and a reverse process that denoises it back toward the original distribution (Ho, Jain, & Abbeel, 2020; Song, Sohl-Dickstein, Kingma, Kumar, Ermon, & Poole, 2021). Score-based variants parameterize the reverse process using a network trained to estimate the score function—the gradient of the log data density—at each noise level. These models have achieved impressive generative quality across image domains, with stable training dynamics that GANs often lack.

DiffPure (Nie, Guo, Huang, Xiao, Vahdat, & Anandkumar, 2022) adapts this for adversarial purification. The idea is straightforward: apply a modest number of forward diffusion steps to a corrupted input, adding enough noise to disrupt adversarial structure without destroying semantic content, then run the reverse process to recover a clean reconstruction. Because adversarial perturbations are high-frequency and structured, they tend to be disrupted faster by the forward process than the coarser features that carry identity information. DiffPure showed strong results against PGD and AutoAttack on standard image benchmarks. Whether those gains translate to fingerprint biometrics—where fine ridge-valley patterns are the identity signal—is an open question that this paper addresses.

2.3 Fingerprint Recognition Architectures

Deep CNNs have largely replaced handcrafted fingerprint features—minutiae points, ridge frequency maps—for recognition tasks (Wang, Zhang, & Zhao, 2022; Cao, & Jain, 2020). EfficientNet (Tan, & Le, 2019) is a well-regarded architecture family that scales depth, width, and resolution jointly using a compound coefficient, achieving strong accuracy-efficiency trade-offs on image benchmarks. It has been applied to fingerprint liveness detection, spoof detection, and cross-sensor matching with good results (Liu, Zhao, Zhang, & Xi, 2012; Toosi, Bottino, Cumani, Negri, & Sottile, 2017). Its parameter efficiency makes it a practical choice for financial authentication systems, where low inference latency and high accuracy are both requirements that cannot be relaxed.

For identity verification specifically, metric learning has become the dominant paradigm: the network is trained to produce embeddings where same-identity images cluster together and different-identity images are far apart (Schroff, Kalenichenko, & Philbin, 2015). Cosine similarity matching, which we used in our prior face recognition work (Ma, & Wilson, 2023), fits naturally into this framework and has been widely adopted in fingerprint systems for its simplicity and robustness to embedding scale variation. ArcFace (Deng, Guo, Xue, & Zafeiriou, 2019), which adds an angular margin penalty during training to sharpen class separation in the embedding space, has become the standard loss function for biometric metric learning.

3. Methodology

3.1 Problem Formulation and Threat Model

Let $f: X \rightarrow Z$ be a fingerprint recognition model mapping an input image x to an embedding z. Verification works by computing cosine similarity sim(z_query, z_gallery) between a query and a registered gallery embedding, accepting the identity if the score exceeds a threshold τ. We use Equal Error Rate (EER)—the operating point where False Acceptance Rate equals False Rejection Rate—as the primary evaluation metric, since it summarizes performance across thresholds without requiring a fixed operating point.

We work under a white-box threat model: the adversary knows the model's parameters and architecture and constructs perturbations δ constrained to an L-infinity ball, $\|\delta\|\infty \leq \varepsilon$. PGD solves the inner maximization: $\delta^* = \text{argmax } \|\delta\|\infty \leq \varepsilon \, L(f(x + \delta), y)$, iterating gradient ascent steps with projection back onto the epsilon-ball. The resulting adversarial input $x\_adv = x + \delta^*$ is evaluated under both evasion (genuine user rejected) and impersonation (impostor accepted) objectives.

3.2 Backbone Model: EfficientNet-B3 for Fingerprint Recognition

We use EfficientNet-B3 as the backbone, initialized with ImageNet pretrained weights and fine-tuned on fingerprint data. The classification head is replaced with a 512-dimensional embedding layer followed by L2 normalization, enabling cosine similarity matching. Training uses ArcFace loss (Deng, Guo, Xue, & Zafeiriou, 2019), which applies an angular margin penalty in the embedding space to sharpen class separation—now standard practice for biometric metric learning.

Fingerprint images are resized to 300×300 pixels, histogram-equalized to normalize intensity across acquisition conditions, and standardized by training-set mean and variance. During standard (non-adversarial) training, augmentations include random horizontal flipping, rotation within ±15 degrees, and Gaussian blur to improve sensor generalization. These augmentations are applied before adversarial perturbation generation during adversarial training, so the model sees realistic variation alongside adversarial noise.

3.3 PGD-Based Adversarial Training

Adversarial training minimizes a robust objective: min_θ E_(x,y)~D [max_||δ||∞≤ε L(f_θ(x + δ), y)], where the inner maximization is solved approximately using K-step PGD with step size α: x_(t+1) = Π_(x + B_ε) (x_t + α · sign(∇_x L(f_θ(x_t), y))). We use ε = 8/255, K = 7, and α = 2/255, the standard PGD-7 configuration (Madry, Makelov, Schmidt, Tsipras, & Vladu, 2018).

To avoid sacrificing too much clean accuracy, we use a mixed training protocol where each mini-batch contains equal proportions of clean and PGD-perturbed samples. Training runs for 50 epochs with Adam, using a cosine annealing schedule from 1×10^{-3} down to 1×10^{-5}, with batch size 32. Perturbations are generated online at each training step rather than precomputed, so the model is exposed to a diverse range of adversarial examples throughout training rather than a fixed set it might memorize.

3.4 Diffusion Model-Based Purification (DiffPure)

The purification module is a DDPM [15] with a UNet backbone (Ronneberger, Fischer, & Brox, 2015), trained on the SOCOFing training images to learn the distribution of clean fingerprints. The forward process adds noise progressively: q(x_t | x_0) = N(x_t; √(ᾱ_t) x_0, (1 - ᾱ_t) I), where ᾱ_t is the cumulative noise schedule. The reverse process is learned as: p_θ(x_(t-1) | x_t) = N(x_(t-1); μ_θ(x_t, t), Σ_θ(x_t, t)).

At inference, given an adversarial image x_adv, we run t* forward steps to obtain a noisy version: x_(t*) = √(ᾱ_(t*)) x_adv + √(1 - ᾱ_(t*)) ε, then apply the reverse process from t* back to 0 to recover x_purified. The t* hyperparameter controls the purification-fidelity trade-off: too few steps and adversarial structure survives; too many and fingerprint ridge detail is lost. We set t* = 100 out of T = 1000 total timesteps, selected empirically on the validation set. To keep inference time manageable, we use the DDIM sampler (Song, Meng, & Ermon, 2021) with S = 50 steps rather than the full DDPM chain.

3.5 Two-Stage Defense Pipeline

At inference time, the pipeline runs as follows: (1) the incoming fingerprint x passes through DiffPure to produce x_purified; (2) x_purified is fed into the adversarially trained EfficientNet-B3 to extract embedding z = f_θ(x_purified); (3) cosine similarity is computed against the registered gallery embedding z_gallery; and (4) the identity is accepted if sim(z, z_gallery) > τ, with τ set at the EER operating point on the clean validation set. The design is intentionally layered: DiffPure handles the bulk of adversarial noise removal, and the adversarially trained backbone handles whatever residual perturbation survives purification. Neither stage is sufficient alone—together, they achieve substantially lower EER than applied independently.

4. Experimental Results and Discussion

4.1 Dataset and Experimental Setup

We use the SOCOFing (Sokoto Coventry Fingerprint) dataset (Shehu, Ruiz-Garcia, Palade, & James, 2018), which contains 6,000 fingerprint images from 600 African subjects—10 images each covering all 10 fingers. The dataset is split by subject: 480 for training (4,800 images), 60 for validation (600 images), and 60 for testing (600 images). For verification evaluation, we sample 1,000 genuine pairs and 5,000 impostor pairs from the test set. Genuine pairs come from different fingers of the same subject; impostor pairs draw one image from each of two different subjects. Adversarial attacks are applied to the query image in each pair.

Table 1. Comparison of EER across attack conditions and defense configurations

Defense Configuration	Clean EER	FGSM (ε=8/255)	PGD-7 (ε=8/255)	PGD-20 (ε=8/255)
No Defense (Baseline)	0.18	0.31	0.42	0.48
DiffPure Only	0.22	0.24	0.29	0.34
Adversarial Training Only	0.21	0.25	0.28	0.33
Proposed (AT + DiffPure)	0.23	0.21	0.19	0.24

Table 1 shows EER across four defense configurations and three attack conditions. The undefended model starts at EER 0.18 on clean inputs—a solid baseline. Under FGSM (ε = 8/255), EER climbs to 0.31. Under PGD-7 and

PGD-20 it reaches 0.42 and 0.48 respectively, confirming that iterative attacks cause substantially more damage than single-step attacks.

DiffPure alone brings PGD-7 EER down from 0.42 to 0.29—a 31% relative reduction—showing that diffusion purification meaningfully strips adversarial noise from fingerprint images. Adversarial training alone achieves PGD-7 EER of 0.28, a similar level of protection. What is more interesting is what happens when both are combined: the proposed AT + DiffPure framework reaches PGD-7 EER of 0.19, a 55% relative improvement over the undefended baseline and a clear step beyond either individual defense. The complementarity works as expected—DiffPure reduces the noise level the backbone must handle, and adversarial training equips the backbone to deal with whatever residual artifacts survive purification.

The clean EER of the combined framework (0.23) is only slightly above the undefended baseline (0.18). The small increase comes from DiffPure's reverse diffusion process smoothing fine ridge detail at $t^* = 100$. This is a modest and acceptable trade-off for the robustness gains achieved. Adaptive t^* scheduling based on input characteristics could reduce this gap further and is worth exploring.

4.2 Ablation Study: Effect of PGD Training Steps and Diffusion Timesteps

Table 2 presents an ablation study examining the effect of the number of PGD training steps $K \in \{3, 7, 10, 20\}$ and the diffusion purification timestep $t^* \in \{50, 100, 150, 200\}$ on PGD-7 test EER. We observe that increasing K from 3 to 7 yields a substantial EER reduction from 0.26 to 0.19, while further increasing K to 20 provides only marginal gains (EER 0.18) at significantly higher training cost. This suggests that PGD-7 adversarial training strikes a favorable trade-off between robustness and efficiency.

Table 2. Ablation study of PGD training steps and diffusion purification timesteps on PGD-7 test EER

PGD Steps (K)	t*=50	t*=100	t*=150	t*=200
K=3	0.28	0.26	0.25	0.27
K=7	0.22	0.19	0.20	0.23
K=10	0.21	0.18	0.19	0.22
K=20	0.21	0.18	0.18	0.21

Across all K values, $t^* = 100$ consistently gives the best EER. At $t^* = 50$, not enough noise is added to disrupt adversarial structure, so perturbations survive purification. At $t^* = 150$, the reverse process smooths out too much ridge detail, raising EER despite stronger noise injection. The non-monotonic pattern highlights a practical design consideration: purification intensity needs to be calibrated to the specific data domain. What works for natural images may not transfer directly to fingerprints, where fine ridge-valley structure carries the identity signal.

4.3 Discussion

The results confirm the core hypothesis: adversarial training and diffusion purification are complementary, and combining them outperforms either applied alone. The 55% relative EER reduction under PGD-7, alongside only a modest 28% relative increase in clean EER, makes a reasonable case for deployment in financial authentication contexts where both security and user experience matter. The trade-off is acceptable—a small clean-performance cost in exchange for substantially stronger protection against iterative attacks.

A few limitations are worth being upfront about. The diffusion purification step adds inference latency. Using DDIM with $S = 50$ steps helps considerably relative to full DDPM sampling, but may still be a bottleneck in high-throughput authentication pipelines where sub-100ms response times are expected. Lightweight diffusion architectures and further sampling acceleration are practical avenues for addressing this. The evaluation also assumes a white-box attacker with full model access—the hardest case for the defender. Black-box and adaptive attack settings, where the adversary knows the purification mechanism is in place, are important to evaluate and remain as future work. Finally, SOCOFing is a relatively small dataset. Results may not generalize uniformly across all sensor types, environmental conditions, and demographic groups found in real-world deployments, and validation on larger corpora like NIST SD14 would strengthen the conclusions.

5. Conclusions

We presented a two-stage adversarial defense for fingerprint authentication in digital financial systems, combining PGD-aware adversarial training with diffusion model-based purification via DiffPure. On the SOCOFing dataset, the framework achieves PGD-7 EER of 0.19—a 55% relative improvement over the undefended baseline of 0.42—while keeping clean EER at 0.23. Neither stage achieves this on its own; the gains come specifically from their combination. Going forward, the most pressing directions are reducing purification latency for real-time deployment, testing under black-box and adaptive attack conditions, and exploring multimodal biometric fusion that pairs fingerprint with face or voice to add further layers of security to digital financial identity systems.

References

Akhtar, N., & Mian, A. (2018). Threat of adversarial attacks on deep learning in computer vision: A survey. *IEEE Access, 6*, 14410-14430.

Bhatt, H. S., Bharadwaj, S., Singh, R., & Vatsa, M. (2015). Recognizing surgically altered face images using multiobjective evolutionary algorithm. *IEEE Transactions on Information Forensics and Security, 8*(1), 89-100.

Cao, K., & Jain, A. K. (2020). Automated latent fingerprint recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 41*(4), 788-800.

Cohen, J., Rosenfeld, E., & Kolter, Z. (2019). *Certified adversarial robustness via randomized smoothing*. ICML.

Delgado-Escano, R., Castro, F. M., Cozar, J. R., Marin-Jimenez, M. J., Guil, N., & de la Torre, F. (2020). An end-to-end multi-task and fusion CNN for inertial-based gait recognition. *IEEE Access, 8*, implement.

Deng, J., Guo, J., Xue, N., & Zafeiriou, S. (2019). *ArcFace: Additive angular margin loss for deep face recognition.* CVPR.

Engelsma, J. J., & Jain, A. K. (2022). Generalizing fingerprint spoof detector: Learning a one-class classifier. *IEEE Transactions on Information Forensics and Security, 17*, 1268-1282.

Engelsma, J. J., Cao, K., & Jain, A. K. (2021). Learning a fixed-length fingerprint representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 43*(6), 1981-1997.

Fang, Z., Czajka, A., & Bowyer, K. W. (2021). Iris presentation attack detection by attention-based and deep pixel-wise binary supervision network. *IEEE Transactions on Information Forensics and Security, 16*, 3791-3803.

Galbally, J., Marcel, S., & Fierrez, J. (2014). Biometric antispoofing methods: A survey in face recognition. *IEEE Access, 2*, 1530-1552.

Ghiani, L., Yambay, D., Mura, V., Tocco, S., Marcialis, G. L., Roli, F., & Schuckers, S. (2013). *LivDet 2013 fingerprint liveness detection competition*. IJCB.

Goodfellow, I. J., Shlens, J., & Szegedy, C. (2015). Explaining and harnessing adversarial examples. ICLR.

Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. *NeurIPS, 33,* 6840-6851.

Jain, A. K., Ross, A., & Prabhakar, S. (2004). An introduction to biometric recognition. *IEEE Transactions on Circuits and Systems for Video Technology, 14*(1), 4-20.

Li, C., Li, J., & Chen, H. (2021). A review of biometric authentication systems for mobile devices. *IEEE Transactions on Information Forensics and Security, 16*(4), 1232-1248.

Liu, F., Zhao, Q., Zhang, D., & Xi, D. (2012). *Fingerprint pore matching based on sparse representation.* ICPR.

Ma, J., & Wilson, A. (2023). A novel domain adaptation-based framework for face recognition under darkened and overexposed situations. *Artificial Intelligence Advances, 5*(1), 63-71.

Ma, J., & Wilson, A. (2023). Mitigating FGSM-based white-box attacks using convolutional autoencoders for face recognition. *Optimizations in Applied Machine Learning, 3*(3).

Madry, A., Makelov, A., Schmidt, L., Tsipras, D., & Vladu, A. (2018). *Towards deep learning models resistant to adversarial attacks*. ICLR.

Maltoni, D., Maio, D., Jain, A. K., & Prabhakar, S. (2009). Handbook of fingerprint recognition. *Springer Science & Business Media.*

Nie, W., Guo, B., Huang, Y., Xiao, C., Vahdat, A., & Anandkumar, A. (2022). *Diffusion models for adversarial purification*. ICML.

Papernot, N., McDaniel, P., Wu, X., Jha, S., & Swami, A. (2016). *Distillation as a defense to adversarial perturbations against deep neural networks.* IEEE S&P.

Ronneberger, O., Fischer, P., & Brox, T. (2015). *U-net: Convolutional networks for biomedical image segmentation*. MICCAI.

Schroff, F., Kalenichenko, D., & Philbin, J. (2015). *FaceNet: A unified embedding for face recognition and clustering*. CVPR.

Shafahi, A., Najibi, M., Ghiasi, A., Xu, Z., Dickerson, J., Studer, C., ... Goldstein, T. (2019). *Adversarial training for free!* NeurIPS.

Shehu, Y. I., Ruiz-Garcia, A., Palade, V., & James, A. (2018). *Sokoto coventry fingerprint dataset (SOCOFing).* arXiv:1807.10609.

Song, J., Meng, C., & Ermon, S. (2021). *Denoising diffusion implicit models.* ICLR.

Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., & Poole, B. (2021). *Score-based generative modeling through stochastic differential equations*. ICLR.

Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., & Fergus, R. (2014). *Intriguing properties of neural networks.* ICLR.

Tan, M., & Le, Q. (2019). *EfficientNet: Rethinking model scaling for convolutional neural networks.* ICML.

Toosi, A., Bottino, A., Cumani, S., Negri, P., & Sottile, P. L. (2017). Feature fusion for fingerprint liveness detection. *IEEE Access, 5*, 23695-23709.

Wang, J., Zhang, L., & Zhao, H. (2022). Deep learning-based fingerprint recognition: A survey. *Pattern Recognition, 119,* 108066.

Xie, C., Wang, J., Zhang, Z., Ren, Z., & Yuille, A. (2018). *Mitigating adversarial effects through randomization.* ICLR.

Yang, W., Wang, S., Hu, J., Zheng, G., & Valli, C. (2019). Security and accuracy of fingerprint-based biometrics: A review. *Symmetry, 11*(2), 141.