Combined Adversarial Attacks for Breaking Fingerprint Recognition and Presentation Attack Detection

Simone Zedda, Simone Maurizio La Cava, Simone Carta, Roberto Casula, Gian Luca Marcialis
Department of Electrical and Electronic Engineering, *University of Cagliari* Cagliari, Italy
email: s.zedda29@studenti.unica.it, {simonem.lac, simone.carta97, roberto.casula, marcialis}@unica.it

***Abstract*—Recent advances in machine and deep learning played a crucial role in developing Fingerprint Presentation Attack Detection (FPAD) modules, aiming at mitigating attacks leveraging artificial fingerprint replicas against Automated Fingerprint Identification Systems (AFISs). However, these advances also introduce a new threat: adversarial attacks designed to mislead the detectors. The risk posed by these attacks is supported by the recent findings on the feasibility of transferring adversarial fingerprint attacks from the digital domain to the physical one. To contribute to the development of countermeasures against this threat, the goal of this work is to investigate the risk posed by state-of-the-art adversarial attacks and strategic combinations between them. Accordingly, we provide an extensive analysis of the impact of white-box and black-box attacks on benchmark FPAD modules in fingerprint recognition scenarios, discussing the most relevant findings for addressing the vulnerability of these detectors to realistic and feasible adversarial attacks.**

***Index Terms*—fingerprint recognition, adversarial perturbation, presentation attack, biometric systems.**

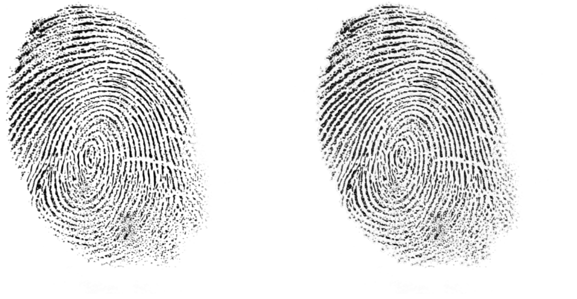
(a) Spoofed fingerprint (b) Perturbed spoofed fingerprint

Fig. 1. Example of spoofed fingerprint (a), classified as fake, and its perturbed version after applying the proposed adversarial attack strategy (b), classified as Live. Specifically, the latter shows the outcome after applying first APGD [15], then DeepFool [16] in the sequential combination (i.e., $ADV_{A,D}$)

I. Introduction

In recent years, we have seen a widespread diffusion of Automated Fingerprint Identification Systems (AFISs) in public security and personal devices, thanks to their reliability and user-friendliness. However, these biometric systems are vulnerable to spoofing attacks, namely, presentation attacks perpetrated by submitting artificial fingerprint replicas to the contact-based or contactless sensor to impersonate an authorized user [1], [2]. Fingerprint Liveness Detection (FLD) systems, also known as Fingerprint Presentation Attack Detection (FPAD) systems, have been developed to counteract these attacks. Typically relying on machine learning and deep learning approaches, these systems have been demonstrated to be accurate in discriminating between genuine and spoofed fingerprints [3].

However, these approaches increase the vulnerability to adversarial attacks, aiming at modifying the classification outcome by digitally perturbing the input image. These attacks, which can be perpetrated at either training or test times (poisoning and evasion, respectively), are especially dangerous due to their capability to modify a classifier's decision while keeping the perturbed image visually unchanged [4]. This vulnerability became even more critical in fingerprint recognition considering the recent findings about the possibility of transferring the deceiving alterations in the physical domain, therefore allowing the creation of Presentation Attack Instruments (PAIs) for attacking AFISs [5]. Hence, this threat may undermine the reliability and, consequently, the trustworthiness of the use of fingerprints for identity recognition, especially considering the robustness required by real-world high-security and forensic applications [6]–[8].

So far, the research community has mainly focused on generating perturbed fake fingerprints that are recognized as bona fide by target FPADs [9]. However, tailoring the adversarial attacks to the deception of the only liveness detection could negatively impact the identity-matching capability of the resulting spoofs due to the introduced large distortions [10]. Therefore, more effort should be dedicated to minimizing adversarial distortions so that perturbed fingerprints can successfully deceive FPADs while maintaining their ability to mislead AFISs.

To reach this goal, the strategic combination of adversarial attacks could allow limiting the perturbation introduced by the individual contributing attacks while still preserving the overall deceiving capability (e.g., Figure 1). This approach, already explored concerning adversarial attacks in other application fields, could represent a feasible trade-off on the effectiveness of misleading both AFISs and FPADs (e.g., [11], [12]).

In this paper, we investigated the effectiveness of single and combined adversarial attacks in a fingerprint recognition scenario. The goal is to evaluate the impact of controlled perturbations that preserve their ability to mislead the AFIS whenever they are integrated into traditional presentation attacks through spoofed fingerprints. Specifically, our contributions are the following: (*i*) a strategy for sequentially combining adversarial attacks in order to improve the effectiveness in deceiving FPADs while preserving the capability of misleading AFISs; (*ii*) extensive analysis of the impact of two state-of-the-art adversarial attacks and their sequential combinations on benchmark FPAD modules, simulating scenarios with various prior knowledge about the attacked system; (*iii*) the assessment of the effect of adversarial perturbations on the matching capability of a baseline AFIS; (*iv*) insights and guidelines to aid the development of countermeasures to novel adversarial attacks.

The rest of the paper is organized as follows. Section II discusses the state of the art of adversarial attacks with the main focus on fingerprint recognition. Section III describes the proposed attack strategy and the preliminary analysis conducted to set the parameters for the individual attacks and their combinations. Section IV presents the experimental framework and Section V discusses the obtained results. The most relevant insights based on results are reported in Section VI. Finally, conclusions are drawn in Section VII.

II. Related Works

Adversarial attacks use subtle input perturbations to exploit model vulnerabilities and cause misclassification. Even imperceptible changes can mislead systems, for instance, causing fingerprint misidentification in biometrics. In this section, we provide a brief introduction to adversarial attacks (Section II-A) and the particular case of their application to fingerprint biometrics (Section II-B).

A. Adversarial Attacks to Image Classification Systems

Adversarial attacks are commonly divided into white box and black box, based on the knowledge of the target model.

White-box attacks rely on full knowledge of the machine learning model, including its architecture, weights, and gradients, enabling the creation of targeted perturbations. Key methods include Fast Gradient Sign Method (FGSM) [13], which alters inputs along the loss gradient; Projected Gradient Descent (PGD) [14], which iteratively refines FGSM-like perturbations; Auto Projected Gradient Descent (APGD) [15], which adapts step sizes for efficiency; DeepFool [16], which minimizes perturbation to cross the classifier's decision boundary; and Carlini & Wagner (C&W) [17], which conceptualizes the attack as an optimization problem to find the smallest perturbation that induces misclassification while maximizing confidence in the incorrect class through the use of a loss function and a confidence parameter.

Black-box attacks lack model access and rely on input-output behavior. A common approach is the transfer attack [18], which employs adversarial examples generated on a white-box model (or surrogate) to fool a black-box target model. Query-based methods [19] involve the attacker repeatedly querying the target classifier with slightly modified inputs, observing the responses to refine perturbations. Another notable technique is the one-pixel attack [21], which aims to fool the classifier by modifying a few key pixels.

B. Adversarial Fingerprint Attacks

Growing security concerns around adversarial attacks led the research community to investigate their potential application in fingerprint recognition.

In the digital domain, adversarial attacks can fool biometric recognition based on convolutional neural networks (CNNs), as demonstrated by attacks on AFIS systems with FGSM, Deep-Fool, and One-Pixel Attack.

These attacks can also be exploited in the physical domain to bypass FPADs. For instance, to our knowledge, Marrone et al. [5] conducted the first adversarial attack on a CNN-based PAD with DeepFool. In particular, they revealed the feasibility of physical adversarial fingerprint presentation attacks (ADV-PAs) by applying liquid latex on perturbed fingerprints, showing that the introduced adversarial perturbations incorporated in the spoofs are maintained during the cast printing process [10]. Specifically, their approach uses a multi-stage ADV-PA with Focus Attention (FA), which applies binarization before and after each attack iteration to enhance perturbation effectiveness.

Variants include Uniform Focus Attention (UFA), distributing perturbations across different fingerprint regions at different iterations, and Robust Focus Attention (RFA), which adds a dilation step to increase robustness and coverage around distinctive fingerprint features.

III. Attack Combination Strategy

As previously introduced, the combination of adversarial attacks can be leveraged to strengthen the degradation of classification accuracy. Therefore, one of our goals is to investigate the potentialities of the sequential combination of attacks against liveness detection in fingerprint recognition scenarios.

Similarly to Casula et al. [10], we considered APGD and DeepFool as adversarial attacks, already introduced in Section II. In particular, we evaluated the impact of attack parameters on the spoofed fingerprints, analyzing the suitable trade-offs between the introduced perturbation and the ability of misleading FPADs.

In the rest of this section, we introduce the FPAD used as the target for white-box attacks (Section III-A) and the approach employed to identify the optimal parameters for individual adversarial attacks (Section III-B).

A. FPAD

To perform the white-box analysis and, therefore, set the parameters of the individual attacks, we employed SimpleCNN [23], a compact CNN including two convolutional blocks and two fully connected layers. Specifically, it is optimized through Adam and early stopping based on binary cross-entropy. The architecture is represented in Figure 2. The images were converted to grayscale and cropped according to the dimensions required by the classifier. Importantly, the

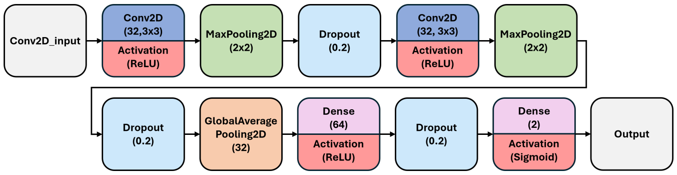

Fig. 2. SimpleCNN architecture [23].

network was trained on the LivDet 2019 dataset to ensure consistency with the test set and maintain coherence with the detectors used during the black-box evaluation phase.

B. Parameter Settings

Regarding both individual adversarial attacks, the effectiveness of the introduced perturbation is influenced by two parameters:

- The *maximum perturbation* allowed for the attack (ϵ), indicating the level of distortion applied to the image. In this case, the study aimed to balance the trade-off between the effectiveness of the attack and the preservation of the distinctive features required for fingerprint matching. Specifically, we analyzed the values from 0.1 to 0.5, where higher values correspond to greater perturbations.
- The *Number of Iterations* (*max_iter*), defining the number of steps the attack algorithm takes to optimize the perturbation, with a higher number leading to a significant increase in the success rate of the attack. Based on the impact of the perturbations introduced by the two adversarial techniques [10], we analyzed *max_iter* ranging from 1 to 3 for APGD and from 30 to 50 for DeepFool.

Starting from the maximum values of the two parameters, all correctly recognized fakes were perturbed using the UFA technique, previously presented in Section II, and with both attack algorithms. In particular, we employed the UFA technique because the digital fingerprint is well-suited for the printing phase and optimized for a subsequent attack in the physical domain.

We tested multiple combinations of parameters through the Coordinate Descent algorithm to determine the maximum parameters for which individual white-box attacks are still ineffective. Specifically, the maximum values for which spoofs are correctly recognized as attacks are the following:

- APGD: $\epsilon = 0.1$, $max\_iter = 1$
- DeepFool: $\epsilon = 0.3$, $max\_iter = 30$

These parameters have been employed by both individual attacks and their combined counterparts. It is important to note that the approach followed for the choice of parameters has been chosen to highlight the potential of the sequential combination. Therefore, the choice of parameters must be modified according to the objective and the requirements of the specific application scenario.

IV. Experimental Framework

Coherently with the aim of this study, we assess the vulnerability of benchmark fingerprint liveness detection systems to a novel attack strategy, sequentially combining two different state-of-the-art adversarial attack techniques while limiting the perturbation on spoofed fingerprints. The rest of this section reports the analyzed data, AFIS, and FPADs.

A. Dataset

We performed the analysis on the LivDet 2019 competition dataset [22]. It comprises fingerprint images of both lives and fakes acquired at 500 dpi from three sensors: GreenBit DactyScan84C, Orcanthus Certis2 Image, and Digital Persona U.are.U 5160. The acquisitions were performed on six fingers per user, i.e., three per hand (thumb, index, and middle). Notably, the spoof test set uses materials different from those used in the training set to simulate "never-seen-before" conditions. The choice of this dataset is also motivated by the availability of pre-tested black-box algorithms, which allowed for a robust evaluation of the proposed adversarial attack. Specifically, we considered data from the first sensor to assess the effectiveness of individual and combined adversarial perturbations.

B. Liveness Detectors

We analyzed the effectiveness of adversarial attacks on different detectors, based on the category of attack, namely white box and black box.

Considering the first type of attack, we employed SimpleCNN, previously introduced in Section III.

To assess the effectiveness of the adversarial attack strategy in a black-box scenario, we considered the liveness detectors that perform the best on images acquired using the DactyScan84C scanner presented in the LivDet 2019 competition [22]. Specifically, we selected JLW and ZJUT as detectors based on the deep learning approach and PADUnkFv as the detector based on hand-crafted features.

C. AFIS

We evaluated the effect of adversarial attacks on matching using the *bozorth3* matcher from the NIST suite[1]. In particular, we verified whether the introduced perturbations compromised the verification performance concerning the ability of the system to match the resulting image with the corresponding genuine fingerprints.

D. Performance Evaluation

We evaluated the impact of both individual and combined adversarial attacks. To provide a detailed investigation, we analyzed the sequential combination between attacks considering all possible application orders, i.e., first APGD then DeepFool ($ADV_{A,D}$) and vice versa ($ADV_{D,A}$).

We assessed the effect of adversarial perturbations on liveness detection performance by analyzing the impact on liveness scores, namely, the probability estimated by the FPAD that the samples are *bona fide* and, therefore, not spoofs. From these scores, we also computed and discussed the Attack Presentation Classification Error Rate (APCER), which measures

[1] https://www.nist.gov/services-resources/software/nist-biometric-image-software-nbis

TABLE I
AVERAGE LIVENESS SCORES ± STANDARD DEVIATION [%] OBTAINED FROM FPADS (THE HIGHEST LIVENESS SCORES FOR EACH FPAD ARE IN BOLD).

Attack	**FPAD**			
	SimpleCNN	**JLW**	**PADUnkFv**	**ZJUT**
Original	4.09±9.32	1.52±5.87	29.25±7.70	1.48±5.73
APGD	0.29±2.91	98.34±9.46	49.25±10.84	98.29±9.46
DeepFool	1.18±6.21	99.43±5.32	48.54±10.76	99.43±5.31
$ADV_{D,A}$	26.14±32.35	99.68±3.82	49.52±10.81	99.68±3.81
$ADV_{A,D}$	**51.28±36.86**	**100.0±0.0**	**49.73±10.72**	**100.0±0.0**

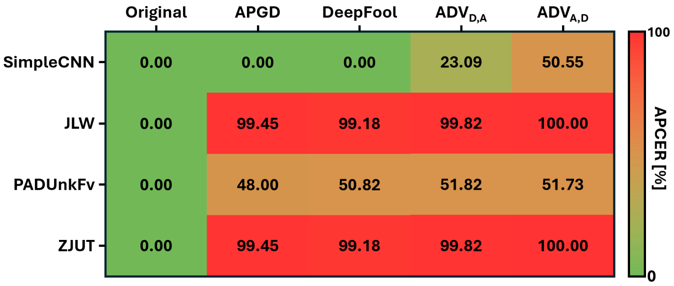

Fig. 3. APCER [%] obtained using 50% as the threshold on liveness scores.

the percentage of fake fingerprints incorrectly classified as *bona fide* by the liveness detection system. Specifically, we obtained the APCER employing a threshold on liveness scores of 50% (i.e., all the samples related to a score below 50% are considered *attacks*, the others as *bona fide*).

Similarly, we examined the effect of adversarial attacks on matching scores generated by AFIS to reveal their influence on fingerprint recognition capability. Specifically, for each finger of any identity, we compared the images related to the perturbed fingerprint with the corresponding genuine images representing the bona fide version of the same fingerprint.

V. RESULTS

A. Liveness Detection

Both individual adversarial attacks demonstrated significant effectiveness in fooling the FPADs, as shown in Table I. The only contrasting result is that of APGD in the white-box attack on SimpleCNN, which reduces liveness scores compared to the unperturbed spoofs (i.e., averages of 0.29% and 4.09%, respectively).

The results also reveal that combining the single attacks can further reduce the overall reliability of FPADs. This outcome is even more relevant when one observes that no single adversarial attack is always the most impactful against all the investigated detectors, suggesting a complementarity between the adversarial perturbations introduced by the two attacks. Specifically, APGD is the most effective attack on PADUnkFv, while DeepFool is the most impactful on the other detectors. This result is not confirmed by the analysis on APCER, which highlighted the opposite trend regarding the most effective attack for specific detectors (Figure 3).

TABLE II
AVERAGE MATCHING SCORES ± STANDARD DEVIATION [%].

Original	**APGD**	**DeepFool**	$ADV_{D,A}$	$ADV_{A,D}$
42.42±27.97	50.24±30.58	50.41±30.35	50.27±30.53	49.96±29.41

Concerning the sequential combinations, the order of the attacks played a crucial role in the success rate of the adversarial examples in white-box attacks, highlighting that $ADV_{A,D}$ is capable of providing significantly higher liveness scores and, therefore, APCER compared to $ADV_{D,A}$. This suggests that applying APGD first, followed by DeepFool to strengthen the perturbation, was more effective in bypassing the FPADs. The trend in black-box attacks is less evident, revealing similar performance against the individual FPADs.

Among the latter, JLW and ZJUT were particularly vulnerable to adversarial attacks. This suggests that these detectors are highly susceptible to adversarial perturbations, even when the attack is transferred from a white-box model. On the other hand, the detector based on hand-crafted features, PADUnkFv, showed more robustness to adversarial perturbation compared to deep learning detectors, reporting 49.73% as the maximum average score (that is, through $ADV_{A,D}$), still significantly increased compared to the original spoofs (29.25%).

B. Fingerprint Matching

The results obtained from the fingerprint recognition analysis, shown in Table II, revealed similar average matching scores and variability between single and combined adversarial attacks. This outcome underscores the minimal effect of the sequential combination compared to the individual attacks included at the recognition level.

The most relevant finding is that the matching scores increased slightly, on average, after incorporating adversarial perturbations (e.g., from 42.42% on the original spoofs to 49.96% after $ADV_{A,D}$). Consequently, beyond improving the capability of the fingerprint to bypass liveness detection, these perturbations also have a slightly positive impact on the likelihood that the matcher recognises a fake fingerprint as belonging to the intended identity. This outcome can be explained by the capability of the enhancement to highlight the minutiae, hence aiding their extraction and the following matching. Moreover, this also suggests that the adversarial perturbations were subtle enough to deceive the liveness detection system without significantly altering the fingerprint's distinctive features required for matching.

VI. KEY INSIGHTS

The trends observed in Section V allow us to draw some important considerations about adversarial attacks and their impact on liveness detection and fingerprint recognition:

- It is important to tailor the attack to the deception of both the fingerprint matcher and the FPAD through spoofed fingerprints rather than only the latter. This means that the introduced adversarial perturbation must be limited to avoid the loss of characteristics required by the recognition step. Therefore, a feasible trade-off between the enhancement in the capability of misleading the liveness detectors and the decay in the effectiveness of spoofs in deceiving the recognition system is required.

- There is no adversarial attack that is always the best one against all FPADs. However, this complementarity could be exploited through a properly designed sequential combination of different attacks.
- The order of attacks in the sequential combination may have a significant impact on the overall deceiving capability. According to the systems analyzed, this observation is particularly relevant to white-box attacks.
- Introducing properly designed perturbations does not necessarily reduce the ability of spoofs to bypass AFIS recognition. On the contrary, introduced perturbations can also emphasize minutiae, thus increasing the matching probability.

VII. Conclusions

In this work, we evaluated the threat of white-box and black-box combined adversarial attacks against Fingerprint Presentation Attack Detectors in recognition scenarios. We proposed an approach for strategically combining adversarial techniques minimizing the introduced distortions to maintain their ability to bypass AFISs, while still misleading FPADs. To validate our proposal, we assessed the impact of state-of-the-art adversarial techniques and their sequential combinations on four benchmark FPADs and a baseline matcher.

Despite the limited techniques analyzed, results show that adversarial attacks tailored to fingerprint recognition can significantly undermine FPAD reliability. Strategically combining these attacks further increases their effectiveness against both fingerprint recognition systems and FPADs, posing a serious risk to fingerprint-based authentication.

Future work should explore more fingerprint sensors, adversarial attacks, and combination strategies, assessing their impact on additional FPADs and recognition systems. The effectiveness of physical replicas with combined perturbations in realistic scenarios, as well as the vulnerability of contactless sensors, should also be investigated.

We hope that the findings of this preliminary study and the insights provided could contribute to the future development of countermeasures against the emerging threat of the combination of spoofing and adversarial attacks.

Acknowledgment

This work was partially supported by project FAIR (PE00000013) under the NRRP MUR program funded by the EU–NGEU (CUP: J23C24000090007).

References

[1] T. Matsumoto, H. Matsumoto, K. Yamada, S. Hoshino. "Impact of artificial" gummy" fingers on fingerprint systems." In Optical security and counterfeit deterrence techniques IV, 2002, https://doi.org/10.1117/12.462719

[2] L. Priesnitz, R. Casula, J. Kolberg, M. Fang, A. Madhu, C. Rathgeb, G. L. Marcialis, N. Damer, C. Busch. "Mobile contactless fingerprint presentation attack detection: generalizability and explainability." IEEE T-BIOM, 2024, DOI: 10.1109/TBIOM.2024.3403770.

[3] M. Micheletto, R. Casula, G. Orrù, S. Carta, S. Concas, S. M. La Cava, J. Fierrez, G. L. Marcialis. "LivDet2023-fingerprint liveness detection competition: advancing generalization." IJCB , 2023, DOI: 10.1109/IJCB57857.2023.10449291.

[4] S. Marrone, C. Sansone. "Adversarial perturbations against fingerprint based authentication systems." 2019 International Conference on Biometrics (ICB), 2019, DOI: 10.1109/ICB45273.2019.8987399.

[5] S. Marrone, R. Casula, G. Orrù, G. L. Marcialis, C. Sansone, "Fingerprint Adversarial Presentation Attack in the Physical Domain", 25th ICPRworkshop, 2020, DOI: 10.1007/978-3-030-68780-9_42.

[6] A. Roy, N. Memon, A. Ross. "Masterprint: Exploring the vulnerability of partial fingerprint-based authentication systems." IEEE TIFS, 2017, DOI: 10.1109/TIFS.2017.2691658.

[7] T. Bollé, E. Casey, M. Jacquet. "The role of evaluations in reaching decisions using automated systems supporting forensic analysis", Forensic Science International: Digital Investigation 34, 2020, https://doi.org/10.1016/j.fsidi.2020.301016.

[8] S. M. La Cava, G. Orrù, M. Drahansky, G. L. Marcialis, F. Roli. "3D face reconstruction: the road to forensics". ACM Computing Surveys, 2023, https://doi.org/10.1145/3625288.

[9] J. Fei, Z. Xia, P. Yu, "Adversarial attacks on fingerprint liveness detection", J Image Video Proc., 2020, DOI: 10.1186/s13640-020-0490-z.

[10] R. Casula, G. Orrù, S. Marrone, U. Gagliardini, G. L. Marcialis, C. Sansone, "Realistic Fingerprint Presentation Attacks Based on an Adversarial Approach", TIFS, 2024, DOI: 10.1109/TIFS.2023.3327663.

[11] X. Mao, Y. Chen, S. Wang, et al. "Composite adversarial attacks", AAAI, 2021, DOI: 10.48550/arXiv.2012.05434.

[12] E. Andrade, I. Sampaio, J. Guérin, J. Viterbo. "Combining Two Adversarial Attacks Against Person Re-Identification Systems", VISAPP, 2023, https://doi.org/10.5220/0011623800003417.

[13] I. Goodfellow, J. Shlens, C. Szegedy. "Explaining and Harnessing Adversarial Examples", 2014, https://doi.org/10.48550/arXiv.1412.6572.

[14] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, A. Vladu, "Towards deep learning models resistant to adversarial attacks", ICLR, 2018, https://doi.org/10.48550/arXiv.1706.06083.

[15] F. Croce, M. Hein, "Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks", 2020, https://doi.org/10.48550/arXiv.2003.01690.

[16] S-M. Moosavi-Dezfooli, A. Fawzi, P. Frossard, "DeepFool: a simple and accurate method to fool deep neural networks", CVPR, 2016, https://doi.org/10.48550/arXiv.1511.04599.

[17] N. Carlini, D. Wagner, "Towards Evaluating the Robustness of Neural Networks", 2017, DOI: 10.1109/SP.2017.49.

[18] N. Papernot, P. McDaniel, I. Goodfellow, "Transferability in Machine Learning: from Phenomena to Black-Box Attacks using Adversarial Samples", 2016, https://doi.org/10.48550/arXiv.1605.07277.

[19] W. Brendel, J. Rauber, M. Bethge, "Decision-Based Adversarial Attacks: Reliable Attacks Against Black-Box Machine Learning Models", 2017, https://doi.org/10.48550/arXiv.1712.04248.

[20] T. Ru, C. Szegedy. (2019). "Simple Black-box Adversarial Attacks", ICML, 2019, https://doi.org/10.48550/arXiv.1905.07121.

[21] J. Su, D. V. Vargas, K. Sakurai, "One Pixel Attack for Fooling Deep Neural Networks", in IEEE Transactions on Evolutionary Computation, 2019, https://doi.org/10.1109/TEVC.2019.2890858.

[22] G. Orrù, R. Casula, P. Tuveri, C. Bazzoni, G. Dessalvi, M. Micheletto, L. Ghiani, G. L. Marcialis, "LivDet in Action - Fingerprint Liveness Detection Competition 2019", ICB, 2019, https://doi.org/10.1109/ICB45273.2019.8987281 .

[23] S. Carta, R. Casula, G. Orrù, M. Micheletto, G. L. Marcialis, "Interpretability of fingerprint presentation attack detection systems: a look at the "representativeness" of samples against never-seen-before attacks", MVA, 2025, DOI: 10.1007/s00138-025-01666-z.