

Cybersecurity and AI: The PRALab Research Experience

Maura Pintor^{1,*}, Giulia Orrú¹, Davide Maiorca¹, Ambra Demontis¹, Luca Demetrio², Gian Luca Marcialis¹, Battista Biggio¹ and Fabio Roli²

¹Pattern Recognition and Applications Laboratory (PRALab), Department of Electrical and Electronic Engineering, University of Cagliari, Italy

²Department of Informatics, Bioengineering, Robotics, and Systems Engineering, University of Genova

Abstract

We present here the main research topics and activities on the design, security, safety, and robustness of machine learning models developed at the Pattern Recognition and Applications Laboratory (PRALab) of the University of Cagliari. Our findings have significantly contributed to identifying and characterizing the vulnerability of such models to adversarial attacks in the context of real-world applications and proposing robust techniques to make these models more reliable in security-critical scenarios.

Keywords

Machine Learning, Adversarial Machine Learning, Biometrics, Cybersecurity

1. Research Group

The Pattern Recognition and Applications (PRA) Laboratory was founded in 1996. The PRALab has been active for more than 20 years at the University of Cagliari. Its mission is to address fundamental issues for the development of future pattern recognition systems in the context of real applications, focused on creating secure systems for security applications, as reflected by our motto:

there is nothing more practical than a good theory, by Kurt Lewin.

Our activities can be categorized into four highly-interdependent lines: (i) development of theories to solve problems of fundamental research, including multiple classifier systems (our original expertise) and adversarial machine learning; (ii) application of these theories to solve practical problems in the research domains of computer vision for video surveillance and ambient intelligence, computer security, biometrics, document and multimedia categorization, and cybersecurity; (iii) testing and validation of the proposed solutions on real-world data (in-vivo experiments); and (iv) development of prototypes and demonstrators, through which the results of basic research are translated into functional products.

The PRALab team, led by the lab director (Prof. Fabio Roli), consists of 1 full and 4 associate professors (Prof. Giorgio Giacinto, Prof. Battista Biggio, Prof. Luca Diodati, Prof. Giorgio Fumera, Prof. Gian Luca Marcialis), 5 assistant professors (Dr. Ambra Demontis, Dr. Davide Maiorca, Dr. Giulia Orrú, Dr. Maura Pintor, Dr. Lorenzo

Putzu), and more than 10 post-doctoral researchers and PhD students. Our team has provided pioneering contributions in the area of AI/ML security, being the first to demonstrate gradient-based evasion [1] (also known as *adversarial examples*) and poisoning attacks [2], and how to mitigate them, playing a leading role in the establishment and advancement of this research field [3].

2. Research Topics

We are among the first to have studied the impact of adversarial machine learning on security applications such as the analysis of malware [1, 4, 5] (Sect. 2.1). We apply robust techniques to improve malware detection on different settings (Sect. 2.2). Furthermore, we investigate techniques to improve the robustness of learning-based systems used for fingerprint, facial and behavioral biometrics, and we organize every two years a challenge to propose and fairly compare newly-developed techniques for reliable fingerprint authentication (Sect. 2.3).

2.1. Machine Learning Security

Recent progress in AI/ML technologies has reported impressive performances in many tasks, paving the way for their use in safety-critical applications like autonomous driving and cybersecurity. Unfortunately, it has been shown that AI/ML technologies are vulnerable to well-crafted attacks performed by skilled attackers. The key to understanding the security properties of AI/ML technologies is to model the threats, simulate the corresponding attacks, and assess the security properties of the system in these scenarios [3]. As shown in Fig. 1, attacks on AI/ML models can be staged both at testing and at training time.

Evasion Attacks. In this scenario, attackers modify the

Ital-IA 2023: 3rd National Conference on Artificial Intelligence, organized by CINI, May 29–31, 2023, Pisa, Italy

*Corresponding author.

✉ maura.pintor@unica.it (M. Pintor)

© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)




Figure 1: Conceptual representation of test time (evasion) and training time (poisoning) attacks. Evasion attacks manipulate the inputs to enforce their misclassification. Poisoning attacks manipulate the training data to cause errors at test time.

input samples of the system to have them misclassified. Our researchers were the first to show that some popular classification algorithms like Support Vector Machines, Neural Networks [1], and feature selection algorithms are vulnerable to this attack [6]. In [7], our researchers have shown how to improve the efficiency of these attacks while obtaining comparable performance. Depending on the considered scenario, the attacks have to be adapted to be sufficiently realistic. For example, our researchers developed evasion attacks to construct malware that evade the target system while preserving all their functionalities [5, 8, 9, 10]. Furthermore, attackers often do not know all the details regarding the target systems. Our researchers have shown that effective attacks can also be developed in this challenging scenario [11, 9, 10, 12].

Unfortunately, these empirical evaluations are often not correctly conducted [3], and this leads to overestimating the robustness of the considered system: the system seems robust only because the attack fails. Our researchers have developed methodologies and debugging tools that can be used to improve current approaches for empirical security evaluations to make them more reliable [13, 14].

Poisoning Attacks. AI/ML technologies are often retrained during the operative phase to consider changes in the data distribution. In this scenario, an attacker can compromise the training data by injecting samples specifically devised to compromise the learning process, making the classifier unable to classify samples at test time correctly. Our researchers have been the first to show that the Support Vector Machines [2], as well as neural networks [15], feature selection methods [16], and clustering algorithms [17] can be compromised by this attack. Furthermore, our researchers have shown that it is possible to find approximate solutions with smaller computational cost and a comparable effectiveness [15, 18].

Other Attacks. We have also been investigating other attacks on AI/ML models, including backdoor poisoning [19] and reprogramming [20].

2.2. Machine Learning for Robust Malware Detection

Machine learning technologies can indeed help in fighting the wide-spreading of malicious software that infect and harm devices of users. Hence, we focus on the development of smarter detectors that better spots threats in the wild. We analyse the impact of *Infection vectors*, non-binary files that carry out additional malicious codes (e.g., PDF and ActionScript files) [21, 22, 23, 24] by leveraging *static analysis* techniques to extract meaningful features that can be useful for classification. Also, we focus on detecting malicious content stored in *Binary Programs* on X86-64 [25] and Android Applications [26, 27, 28, 29], by investigating how specific API calls (e.g., system-based and crypto-related) can be used as features to discriminate between malicious and benign programs.

Our research also focused on understanding the *robustness* of such detectors against *test-time* evasion attacks. We were among the first to test both *white-box* and *black-box* attacks by constructing *working* samples that reflected the attacker’s modifications. Our research covers the attacking strategies and the manipulations, addressing four major directions.

Robustness of Windows Malware Detectors. We study the manipulations of Windows programs [30, 9, 10, 31], and we develop attacks in both white-box and black-box settings. Also, these strategies have been also used to test the robustness of commercial products hosted on VirusTotal,¹ highlighting that many of them are susceptible to adversarial perturbations as well.

Robustness of Android Malware Detectors. We investigated how Android applications can be modified to inject adversarial content inside them, by adding, e.g., fake permissions and API calls. Our studies focus on state-of-the-art detectors, and we apply our findings against them [5, 32].

Robustness of PDF Malware Detectors. We studied the PDF file format and the practical evasive manipulations that allow keeping all the relevant information of the original file [4, 33, 23], in both white-box and black-box settings.

Mitigation of Adversarial Attacks. After having investigated the weaknesses of machine learning models, we devised novel techniques that are robust to these kinds of perturbations. Specifically, we focus on forcing the attackers to drastically increase their effort to

¹<https://virustotal.com>

bypass detection [5]. When unconstrained, trained linear models tend to rely only upon a few discriminating features to make their predictions. This enables the attacker to easily bypass detection by perturbing only such few, highly-relevant features. To mitigate this issue and improve robustness, we introduced a theoretically-sound regularization term that provides the optimal, robust linear model against such attacks. This classifier practically works by forcing the optimizer to redistribute the importance of many input features, bounding the maximum absolute weight value assigned to each of them. This in turn constrains the attacker to manipulate more features to bypass detection, effectively improving robustness.

Explainable Malware Detection. The focus of this research direction is twofold. First, we include techniques that can be easily explained and verified inside the development of malware detectors, allowing us to understand why a model flags a program as malicious [34, 5]. Second, we dissect the reason why adversarial attacks succeed against the model under test [30], by leveraging explainability techniques proposed in the state of the art.

2.3. Biometrics

One of our main activities in the field of biometrics, is the design of methods and models for detecting attacks that compromise the authenticity of personal identity, particularly fingerprint and facial identity. Indeed, authentication systems are vulnerable to the submission of artificial replicas of the biometric trait on the sensor, known as presentation attacks (PAs). Our researchers, in addition to evaluating the dangers of new manufacturing techniques for PAs from latent fingerprints [35] or through adversarial techniques [36], are involved in the development of appropriate detectors, and their integration into current fingerprint authentication systems [37]. Since 2009, we are organizers of the International Fingerprint Liveness Detection Competition (LivDet)². LivDet is a biennial appointment for companies and research institutions with the aim of assessing the performance of state-of-the-art fingerprint PA detection systems.

Moreover, we study and implement new techniques for deepfake detection, i.e. methods to detect manipulations of facial identity obtained through deep learning techniques in video content, [38].

In addition to the study of recognition systems through strong biometrics, our researchers are specialized in behavioural biometrics. In particular, through the study of natural movements of people such as the speed of creation or disintegration of groups of individuals [39], it is possible to detect the emergence of anomalies, such as episodes of violence or panic.

Among the numerous challenges that our researchers face, we want to highlight the explainable AI for biometrics topic, which is an aspect of fundamental importance for designing reliable and understandable recognition and classification systems, especially from the point of view of forensic analysis.

3. Projects

Our research activities are carried out in the framework of regional, national, and European projects funded by public and private initiatives. We had more than twenty-five projects founded between 2012 and 2020. The full list is available at <http://pralab.diee.unica.it/en/Projects>. Six of them were founded by the European Commission, and two of them were coordinated by the PRALab. Overall, we received 3 million euros of funding, whose the European Commission provided half. The annual turnover is around four hundred thousand euros.

We have different ongoing projects on AI security:

1. 2023-2026 - Sec4AI4Sec aims to devise the testing and protection of AI-enabled components in software security assets. The project will start in late 2023.
2. 2022-2026 - "ELSA: European Lighthouse on Secure and Safe AI," funded by the EU Horizon Europe research and innovation programme (grant no. 101070617).
3. 2020-2023 - FFG COMET Module S3AI: "Security and Safety for Shared Artificial Intelligence," funded by BMK, BMDW, and the Province of Upper Austria in the frame of the COMET Programme managed by the Austrian Research Promotion Agency FFG. This project aims to provide the foundations required to build secure and safe shared artificial intelligence systems.
4. 2019-2023 - PRIN 2017 BullyBuster, funded by the Italian Ministry of Education, University and Research (CUP: F74I19000370001). The project aims to provide AI-based solutions against the phenomenon of bullying and cyberbullying.
5. 2020-2022 - PRIN 2017 RexLearn: "Reliable and Explainable Machine Learning," funded by the Italian Ministry of Education, University and Research (grant no.2017TWNMH2). This project aims to develop novel learning paradigms, able to take reliable and explainable decisions, and to assess and mitigate the security risks associated with potential misuses of machine learning.

Some other relevant projects are listed in the following:

- 2017-2019 - Research and Innovation Action LETS-CROWD: "Law Enforcement agencies human factor methods and Toolkit for the Secu-

²<http://livdet.diee.unica.it>

urity and protection of CROWDs in mass gatherings”. Call: H2020 - SEC-07-FCT-2016-2017. Grant Agreement H2020/N.740466.

- 2015-2018 – Innovation Action DOGANA: “advanced sOcial enGineering And vulNerability Assessment Framework”. Call: H2020 – DS 2014-1. Grant Agreement H2020/N.653618.
- 2014-2016 – CSA CyberROAD: “Development of the Cybercrime and Cyberterrorism Research Roadmap”. Call: FP7 – SEC 2013.2.5-1. Grant Agreement FP7-SEC-2013/N.607642.
- 2014-2016 - ILLBuster, “Buster of ILLegal Contents spread by malicious computer networks”. DGHOMe - ISEC, Prevention of and Fight Against Crime. Grant Agreement: HOME/2012/ISEC/AG/4000004360.
- 2013-2015 - MAVEN: “Management and Authenticity Verification of Multimedia Contents”. Call: FP7-SME- 2013-1. Grant Agreement FP7-SME-2013-1/N.606058.

4. Developed Tools

Machine Learning Security and Robust Malware Detection. As explained in the previous section, correctly evaluating the robustness of AI/ML technologies might be challenging. Our researchers have developed different tools that help to perform security evaluation³. These tools include SecML [40], a Python library to assess the security evaluation of AI/ML technologies against evasion and poisoning attacks, and an extension of this library, called SecML Malware [31] ad-hoc for Windows malware. For each of them, they have released a tool that evaluates security through a graphical interface: PandaVision, and ToucanStrike. Furthermore, our researchers have released a tool to evaluate if an attack is effective in the considered scenario⁴.

Biometrics. In all research fields and in the technology transfer projects they lead, the researchers of the biometrics unit provided proofs-of-concept and tools. For example, within two projects funded by the Italian Presidency of Minister, Scientific Division, our researchers have developed the Fingerprint Forensic tool and the Deepfake detection tool. The first is a tool of advanced fingerprint image processing techniques, including detecting fingerprint PA coming from latent marks. The second is a tool that combines different state-of-the-art deepfake detection algorithms to exploit complementary information in assessing the authenticity of multimedia content.

³<https://github.com/pralab>

⁴<https://github.com/pralab/IndicatorsOfAttackFailure>

5. Challenges and Perspectives

While research is quickly progressing in AI/ML Security, companies are working on automating the development and operations of ML models (MLOps) without focusing too much on ML security-related issues. In this respect, a relevant challenge for the future will be to extend the current MLOps paradigm and also to encompass ML security (towards implementing what we refer to as MLSecOps). To this end, we plan to incorporate research on security testing, protection, and monitoring of AI/ML models into the MLOps development cycle. In particular, we plan to extend our research towards: (i) developing and improving attacks (including evasion, poisoning, and privacy threats) for making security testing and validation of AI/ML models more efficient and available for a wider set of application domains; (ii) designing improved defenses with robustness guarantees to protect AI/ML models not only against such attacks but also to enable reliable classification when out-of-distribution data is provided as input; and (iii) designing methods that constantly monitor if a deployed model is under attack during operation, enabling prompt reaction when needed. We firmly believe that integrating these dimensions into an MLSecOps cycle will definitely help software engineers and developers to seamlessly deploy and maintain more secure, reliable, and trustworthy AI/ML models in practice.

References

- [1] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrndić, P. Laskov, G. Giacinto, F. Roli, Evasion attacks against machine learning at test time, in: ECML PKDD, Part III, volume 8190 of *LNCS*, 2013, pp. 387–402.
- [2] B. Biggio, B. Nelson, P. Laskov, Poisoning attacks against support vector machines, in: 29th ICML, 2012, pp. 1807–1814.
- [3] B. Biggio, F. Roli, Wild patterns: Ten years after the rise of adversarial machine learning, *Patr. Rec.* 84 (2018) 317–331.
- [4] D. Maiorca, I. Corona, G. Giacinto, Looking at the bag is not enough to find the bomb: an evasion of structural methods for malicious pdf files detection, in: Proceedings of the 8th ACM SIGSAC symposium on Information, computer and communications security, ASIA CCS '13, ACM, New York, NY, USA, 2013, pp. 119–130.
- [5] A. Demontis, M. Melis, B. Biggio, D. Maiorca, D. Arp, K. Rieck, I. Corona, G. Giacinto, F. Roli, Yes, machine learning can be more secure! a case study on android malware detection, *IEEE Trans. Dependable and Secure Computing* (2019).
- [6] F. Zhang, P. Chan, B. Biggio, D. Yeung, F. Roli, Ad-

- versarial feature selection against evasion attacks, *IEEE Trans. on Cybernetics* 46 (2016) 766–777.
- [7] M. Pintor, F. Roli, W. Brendel, B. Biggio, Fast Minimum-norm Adversarial Attacks through Adaptive Norm Constraints, in: A. Beygelzimer, Y. Dauphin, P. Liang, J. W. Vaughan (Eds.), *Advances in Neural Information Processing Systems*, 2021.
- [8] B. Kolosnjaji, A. Demontis, B. Biggio, D. Maiorca, G. Giacinto, C. Eckert, F. Roli, Adversarial Malware Binaries: Evading Deep Learning for Malware Detection in Executables, *ArXiv* (2018). [arXiv:1803.04173](https://arxiv.org/abs/1803.04173).
- [9] L. Demetrio, B. Biggio, G. Lagorio, F. Roli, A. Armando, Functionality-preserving black-box optimization of adversarial windows malware, *IEEE Transactions on Information Forensics and Security* 16 (2021) 3469–3478.
- [10] L. Demetrio, S. E. Coull, B. Biggio, G. Lagorio, A. Armando, F. Roli, Adversarial EXEmples: A survey and experimental evaluation of practical attacks on machine learning for windows malware detection, *ACM Trans. Priv. Secur.* 24 (2021).
- [11] A. Demontis, M. Melis, M. Pintor, M. Jagielski, B. Biggio, A. Oprea, C. Nita-Rotaru, F. Roli, Why do adversarial attacks transfer? Explaining transferability of evasion and poisoning attacks, in: *USENIX Security*, USENIX Association, 2019.
- [12] M. Pintor, D. Angioni, A. Sotgiu, L. Demetrio, A. Demontis, B. Biggio, F. Roli, Imagenet-patch: A dataset for benchmarking machine learning robustness against adversarial patches, *Pattern Recognition* 134 (2023) 109064.
- [13] M. Pintor, L. Demetrio, A. Sotgiu, A. Demontis, N. Carlini, B. Biggio, F. Roli, Indicators of Attack Failure: Debugging and Improving Optimization of Adversarial Examples, in: A. H. Oh, A. Agarwal, D. Belgrave, K. Cho (Eds.), *Advances in Neural Information Processing Systems*, 2022. URL: <https://openreview.net/forum?id=Y1sWzKW0k4L>.
- [14] M. Pintor, L. Demetrio, G. Manca, B. Biggio, F. Roli, Slope: A First-order Approach for Measuring Gradient Obfuscation, in: *Proc. of the ESANN, ESANN 2021*, 2021.
- [15] L. Muñoz-González, B. Biggio, A. Demontis, A. Paudice, V. Wongrassamee, E. C. Lupu, F. Roli, Towards poisoning of deep learning algorithms with back-gradient optimization, in: *Proc. of the 10th ACM Works. AISec@CCS 2017*, 2017, pp. 27–38. URL: <https://doi.org/10.1145/3128572.3140451>.
- [16] H. Xiao, B. Biggio, G. Brown, G. Fumera, C. Eckert, F. Roli, Is feature selection secure against training data poisoning?, in: *ICML*, 2015, pp. 1689–1698.
- [17] B. Biggio, I. Pillai, S. R. Bulò, D. Ariu, M. Pelillo, F. Roli, Is data clustering in adversarial settings secure?, in: *Proc. of the 2013 AISec, AISec '13*, New York, NY, USA, 2013, pp. 87–98.
- [18] A. E. Cinà, S. Vascon, A. Demontis, B. Biggio, F. Roli, M. Pelillo, The Hammer and the Nut: Is Bilevel Optimization Really Needed to Poison Linear Classifiers?, in: *IJCNN 2021, Shenzhen, China, July 18-22, 2021*, IEEE, 2021, pp. 1–8. URL: <https://doi.org/10.1109/IJCNN52387.2021.9533557>. doi:10.1109/IJCNN52387.2021.9533557.
- [19] A. E. Cinà, K. Grosse, S. Vascon, A. Demontis, B. Biggio, F. Roli, M. Pelillo, Backdoor learning curves: Explaining backdoor poisoning beyond influence functions, 2021. [arXiv:2106.07214](https://arxiv.org/abs/2106.07214).
- [20] Y. Zheng, X. Feng, Z. Xia, X. Jiang, A. Demontis, M. Pintor, B. Biggio, F. Roli, Why adversarial reprogramming works, when it fails, and how to tell the difference, *Information Sciences* (2023).
- [21] D. Maiorca, G. Giacinto, I. Corona, A pattern recognition system for malicious pdf files detection, in: P. Perner (Ed.), *Machine Learning and Data Mining in Pattern Recognition*, volume 7376 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, 2012, pp. 510–524.
- [22] D. Maiorca, D. Ariu, I. Corona, G. Giacinto, A structural and content-based approach for a precise and robust detection of malicious PDF files, in: O. Camp, E. R. Weippl, C. Bidan, E. Aïmeur (Eds.), *ICISSP 2015 - Proceedings of the 1st International Conference on Information Systems Security and Privacy*, ESEO, Angers, Loire Valley, France, 9-11 February, 2015, SciTePress, 2015, pp. 27–36. URL: <https://doi.org/10.5220/0005264400270036>. doi:10.5220/0005264400270036.
- [23] D. Maiorca, B. Biggio, G. Giacinto, Towards adversarial malware detection: Lessons learned from pdf-based attacks, *ACM Comput. Surv.* 52 (2019) 78:1–78:36. URL: <http://doi.acm.org/10.1145/3332184>. doi:10.1145/3332184.
- [24] D. Maiorca, A. Demontis, B. Biggio, F. Roli, G. Giacinto, Adversarial Detection of Flash Malware: Limitations and Open Issues, *Computers & Security* 96 (2020). URL: https://www.sciencedirect.com/science/article/pii/S0167404820301760?dgcid=rss_sd_all. doi:<https://doi.org/10.1016/j.cose.2020.101901>.

- [25] F. Meloni, A. Sanna, D. Maiorca, G. Giacinto, Extended abstract: Effective call graph fingerprinting for the analysis and classification of windows malware, in: L. Cavallaro, D. Gruss, G. Pellegrino, G. Giacinto (Eds.), *Detection of Intrusions and Malware, and Vulnerability Assessment - 19th International Conference, DIMVA 2022, Cagliari, Italy, June 29 - July 1, 2022, Proceedings*, volume 13358 of *Lecture Notes in Computer Science*, Springer, 2022, pp. 42–52. URL: https://doi.org/10.1007/978-3-031-09484-2_3. doi:10.1007/978-3-031-09484-2_3.
- [26] D. Maiorca, D. Ariu, I. Corona, M. Aresu, G. Giacinto, Stealth attacks: An extended insight into the obfuscation effects on android malware, *Comput. Secur.* 51 (2015) 16–31.
- [27] D. Maiorca, F. Mercaldo, G. Giacinto, C. A. Visaggio, F. Martinelli, R-packdroid: API package-based characterization and detection of mobile ransomware, in: A. Seffah, B. Penzenstadler, C. Alves, X. Peng (Eds.), *Proceedings of the Symposium on Applied Computing, SAC 2017, Marrakech, Morocco, April 3-7, 2017, ACM, 2017*, pp. 1718–1723. URL: <https://doi.org/10.1145/3019612.3019793>. doi:10.1145/3019612.3019793.
- [28] M. Scalas, D. Maiorca, F. Mercaldo, C. A. Visaggio, F. Martinelli, G. Giacinto, On the effectiveness of system api-related information for android ransomware detection, *Comput. Secur.* 86 (2019) 168–182. URL: <https://doi.org/10.1016/j.cose.2019.06.004>. doi:10.1016/j.cose.2019.06.004.
- [29] A. Janovsky, D. Maiorca, D. Macko, V. Matyas, G. Giacinto, A longitudinal study of cryptographic API: A decade of android malware, in: S. D. C. di Vimercati, P. Samarati (Eds.), *Proceedings of the 19th International Conference on Security and Cryptography, SECRYPT 2022, Lisbon, Portugal, July 11-13, 2022, SCITEPRESS, 2022*, pp. 121–133. URL: <https://doi.org/10.5220/0011265300003283>. doi:10.5220/0011265300003283.
- [30] L. Demetrio, B. Biggio, G. Lagorio, F. Roli, A. Armando, Explaining vulnerabilities of deep learning to adversarial malware binaries, in: *Proceedings of the Third Italian Conference on Cyber Security (ITASEC), 2019*.
- [31] L. Demetrio, B. Biggio, Secml-malware: Pentesting windows malware classifiers with adversarial examples in python, *arXiv preprint arXiv:2104.12848* (2021).
- [32] F. Cara, M. Scalas, G. Giacinto, D. Maiorca, On the feasibility of adversarial sample creation using the android system API, *Information* 11 (2020) 433. URL: <https://doi.org/10.3390/info11090433>. doi:10.3390/info11090433.
- [33] D. Maiorca, B. Biggio, Digital investigation of pdf files: Unveiling traces of embedded malware, *IEEE Security Privacy* 17 (2019) 63–71. doi:10.1109/MSEC.2018.2875879.
- [34] M. Melis, M. Scalas, A. Demontis, D. Maiorca, B. Biggio, G. Giacinto, F. Roli, Do gradient-based explanations tell anything about adversarial robustness to android malware?, *International Journal of Machine Learning and Cybernetics* 13 (2022) 217–232.
- [35] R. Casula, M. Micheletto, G. Orrù, G. L. Marcialis, F. Roli, Towards realistic fingerprint presentation attacks: The screen-spoof method, *Pattern Recognition Letters* (2022). URL: <https://www.sciencedirect.com/science/article/pii/S0167865522002653>. doi:<https://doi.org/10.1016/j.patrec.2022.09.002>.
- [36] S. Marrone, R. Casula, G. Orrù, G. Marcialis, C. Sansone, Fingerprint adversarial presentation attack in the physical domain, in: A. Del Bimbo, R. Cucchiara, S. Sclaroff, G. Farinella, T. Mei, M. Bertini, H. Escalante, R. Vezzani (Eds.), *Pattern Recognition. ICPR International Workshops and Challenges, Springer International Publishing, Cham, 2021*, pp. 530–543.
- [37] M. Micheletto, G. L. Marcialis, G. Orrù, F. Roli, Fingerprint recognition with embedded presentation attacks detection: Are we ready?, *IEEE Transactions on Information Forensics and Security* 16 (2021) 5338–5351. doi:10.1109/TIFS.2021.3121201.
- [38] S. Concas, G. Perelli, G. L. Marcialis, G. Puglisi, Tensor-based deepfake detection in scaled and compressed images, in: *2022 IEEE International Conference on Image Processing (ICIP), IEEE, 2022*, pp. 3121–3125.
- [39] G. Orrù, D. Ghiani, M. Pintor, G. L. Marcialis, F. Roli, Detecting anomalies from video-sequences: a novel descriptor, in: *2020 25th International Conference on Pattern Recognition (ICPR), IEEE, 2021*, pp. 4642–4649.
- [40] M. Pintor, L. Demetrio, A. Sotgiu, M. Melis, A. Demontis, B. Biggio, secml: Secure and explainable machine learning in python, *SoftwareX* 18 (2022) 101095.