Maura Pintor, Assistant Professor @ Unica

Qualification	PhD in Electronic and Computer Engineering
Day and Place of Birth	October 20th, 1991, Cagliari (CA), Italy
Nationality	Italian
Affiliation	Department of Electrical and Electronic Engineering (DIEE), University of Cagliari, Piazza d'Armi, 09123, Cagliari Italy

maura.pintor@unica.it
https://maurapintor.github.io
@maurapintor
https://www.linkedin.com/in/maura-pintor/

Education and Research

- 03/2023 - ongoing **University of Cagliari (Italy), Assistant Professor (RTDa).** Machine learning security.
- 10/2021 - 02/2023 **University of Cagliari (Italy), Postdoctoral Researcher.** Machine learning security.
- 2018 - 2022 **University of Cagliari (Italy) - PhD (honors) in Electronic and Computer Engineering**
  Topic: Adversarial Machine Learning.
  Graduation date: 18/02/2022
  Thesis: Towards Debugging and Improving Adversarial Robustness Evaluations.
- 2016 - 2018 **University of Cagliari (Italy) - Telecommunications Engineering, 1st Level Degree (Master).**
  Graduation date: 25/09/2018. Final degree mark: 110/110, magna cum Laude
  Thesis: A novel temporal descriptor for analyzing small and large crowds by computer vision algorithms.
- 2010 - 2016 **University of Cagliari (Italy) - Electronic Engineering, 2nd Level Degree (Bachelor).** Graduation date: 22/07/2016. Final degree mark: 104/110
  Thesis: Methods and Algorithms for gender classification through face image acquisition.

Visiting Researcher

- 06/2024 - 10/2024 **Universitat Autònoma de Barcelona (Spain), Visiting Researcher.** Laboratory: CVC.
- 05/2021 - 08/2021 **Software Competence Center Hagenberg (Austria), Visiting Student.** Laboratory: SCCH.
- 03/2020 - 06/2020 **University of Tübingen (Germany), Visiting Student.** Laboratory: Bethgelab.

Teaching

Invited Lectures, Talks and Summer Schools

- 2026 Lecturer at the 6th IAPR TC10/TC11 Summer School on Next-Gen Document Understanding: RAG, VLMs, and Structured Knowledge Extraction.
- 2025 Lecturer at the First Summer School on AI and Cybersecurity at TU Wien.
- Keynote Speaker at the 7th workshop of ML for Cybersecurity (co-located with ECML-PKDD).
- "Foundations of Machine Learning and its Application to Cybersecurity" at the Summer School on Security and Privacy in the Age of AI, organized by the DistriNet Research Unit at KU Leuven.
- Keynote speaker at the 2025 Alan Turing Women in AI Security Workshop.
- 2024 "Introductory lecture on Security of AI (AML)" (co-lecturer with Prof. Fabio Roli) at the Summer School on Security and Privacy in the Age of AI, organized by the DistriNet Research Unit at KU Leuven.
- "Reliable Evaluation and Benchmarking of Machine Learning Models" at the UPM Cybersecurity Postgraduate Summer School, organized by the Cátedra CiberSeguridad at the Universidad Politécnica de Madrid.
- Keynote speaker at the 3rd Workshop on Rethinking Malware Analysis (WoRMA), co-located with IEEE EuroS&P 24.
- Presented "Machine Learning Security: Lessons Learned and Future Challenges" at the Università della Svizzera Italiana (USI), Lugano, Switzerland.
- Presented "Where ML security is broken and how to fix it" at the Tübingen AI Center, Eberhard Karls University of Tübingen in cooperation with the Max Planck Institute for Intelligent Systems, Germany.
- 2023 "Foundations of Machine Learning and its Application to Cybersecurity" (co-lecturer with Luca Demetrio) at the Summer School on Security and Privacy in the Age of AI, organized by the DistriNet Research Unit at KU Leuven.

Courses

- 2025 - yearly **University of Cagliari (Italy), Course Instructor.** BSc Course on Web Programming, 70 hours.

Teaching (continued)

2024 - yearly **University of Cagliari (Italy), Course Instructor.** PhD Course on Deep Learning and Computer Vision with PyTorch, 20 hours..

2022 - ongoing **University of Cagliari (Italy), Teaching Assistant.** Machine Learning Security (MSc in Computer Engineering, Cybersecurity and Artificial Intelligence).

2019 - ongoing **University of Cagliari (Italy), Teaching Assistant.** Machine Learning (MSc in Computer Engineering, Cybersecurity and Artificial Intelligence).

2019 - 12/2023 **University of Cagliari (Italy), Teaching Assistant.** Industrial Software Development (MSc in Computer Engineering, Cybersecurity and Artificial Intelligence).

Thesis Supervisor (MSc Course in Computer Engineering, Cybersecurity and Artificial Intelligence)

2026 Supervision of the student Marco Pintore for his MSc Thesis "Counterfeit Answers: Adversarial Forgery against OCR-Free Document Visual Question Answering".

Supervision of the student Stefania Macis for her MSc Thesis "Asynchronous Web-Based Security Evaluation of Machine Learning Models".

2023 Co-supervision of the student Luca Scionis for his MSc Thesis "Neural Network Pruning for Adversarial Robustness: An Overview and Experimental Analysis".

Co-supervision of the student Giuseppe Floris for his MSc Thesis "Hyperparameter Optimization for Fast Minimum-norm Attacks".

2021 Co-supervision of the student Giovanni Manca for his MSc Thesis "Understanding Failures of Gradient-based Attacks on Machine Learning".

Co-supervision of the student Giorgio Piras for his MSc Thesis "On Explainability of Machine Learning DGA detectors from DNS traffic data".

Research Projects

11/2024–ongoing Participation, with the University of Cagliari, in the EU project "A Comprehensive Trustworthy Framework for Connected Machine Learning and Secure Interconnected AI Solutions" (CoEvolution), GA no.: 101168560, funded by the EU in the programme HORIZON-CL3-2023-CS-01-03.

10/2023–ongoing Participation, with the University of Cagliari, in the EU project "Security for AI and AI for Security" (Sec4AI4Sec), GA no.: 101120393, funded by the EU in the programme HORIZON-CL3-2022-CS-01.

10/2022–ongoing Participation, with the University of Cagliari, in the EU project "European Lighthouse on Secure and Safe AI" (ELSA), GA no.: 101070617, funded by the EU in the programme HORIZON-CL4-2021-HUMAN-01.

10/2021–04/2023 Participation, with the University of Cagliari, in the research project "Huawei R&D Agreement: Deep Reinforcement Learning Key Security Technologies", GA n. TC20201118006.

03/2021–10/2023 Scientific Coordinator, with the company Pluribus One, of the WP6 (Impact: Benchmark Datasets and Tool Flow Pilots) of the EU project "Assurance and certification in secure Multi-party Open Software and Services" (AssureMOSS), GA no.: 952647, funded by the EU in the programme H2020-SU-ICT-2019.

03/2019–03/2020 Scientific Coordinator, with the company Pluribus One, in the EU project "Software framework for runtime-Adaptive and secure deep Learning On Heterogeneous Architectures" (ALOHA), GA no.: 780788, funded by the EU in the programme H2020-ICT-2017-1.

Employment History

03/2021–10/2023 **Pluribus One S.r.l. (Italy), Collaborator.** Automated techniques to assess, manage, and re-certify the security and privacy risks of multi-party open software and services (MOSS). Project AssureMOSS - EU.

03/2019–03/2020 **Pluribus One S.r.l. (Italy), Collaborator.** Deep Learning systems in low-power heterogeneous platforms. Development of a module for adversarial robustness evaluations. Project ALOHA - EU.

02/2018–07/2018 **Pluribus One S.r.l. (Italy), Software developer.** Systems for Internet traffic security.

Participation to International Groups and Associations and Service

Participation to International Groups and Associations

2024 Member of the European Laboratory for Learning and Intelligent Systems (ELLIS).

2024–current Member of the IEEE Information Forensics and Security Technical Committee (IFS TC).

2018–current Member of the Institute of Electrical and Electronics Engineers (IEEE)

Member of the Association for Computing Machinery (ACM)

Participation to International Groups and Associations and Service (continued)

- Member of the Pattern Recognition and Applications (PRA) Laboratory of the University of Cagliari (Italy).

Workshop and Conference Organization

- 2026 Area Chair for the IEEE International Workshop on Information Forensics and Security (WIFS).
- 2025–2026 Workshop co-chair at International Workshop on Designing and Measuring Security in Systems with AI, co-loc. with IEEE EuroS&P.
- 2024–2025 Track co-chair at Safe, Secure and Robust AI Track at the ACM SAC'24 and '25 Conference.
- 2023–2025 Workshop co-chair at 16th, 17th and 18th ACM Workshop on Art. Int. and Security (AISec '23, '24, and '25), co-loc. with ACM CCS.
- 2024–2026 Area Chair for Neural Information Processing Systems (NeurIPS 2024).

Associate Editor for Journals

- 2025–ongoing CAE for IEEE Transactions on Information Forensics and Security (IEEE T-IFS).
- 2024–ongoing AE for Pattern Recognition.
- 2022–2025 AE for the International Journal of Machine Learning and Cybernetics (IJMLC).

Reviewer for Journals

- IEEE Transactions on Information Forensics and Security (IEEE T-IFS), IEEE Transactions on Image Processing (IEEE TIP), IEEE Transactions on Dependable and Secure Computing (IEEE TDSC), IEEE Transactions on Neural Networks and Learning Systems (IEEE-TNNLS), ACM Transactions on Privacy and Security (TOPS).

Reviewer for Conferences

- 2025 IEEE S&P, USENIX Security, IEEE SaTML, ICLR, IEEE ICCV, IEEE/CVF CVPR, ACM CCS.
- 2024 ACSAC, ICPR, USENIX Security, ACM CCS, ECCV, ICLR (**Top Reviewer**)
- 2023 NeurIPS (**Top Reviewer**), ACSAC, ICCV.
- 2020–ongoing Reviewer for more than 20 workshops co-located with top conferences among the ones listed above.

Research Publications

Below is a list of selected publications. The full and updated list of publications is available at `https://scholar.google.it/citations?user=Tu45bY4AAAAJ&hl=it`.

Top Journal Papers

1. Gupta, S., Angioni, D., **Pintor**, **M.**, Demontis, A., Schönherr, L., Roli, F., & Biggio, B. (2026). Buffer-free class-incremental learning with out-of-distribution detection. Pattern Recognition, 172. https://doi.org/10.1016/j.patcog.2025.112441

2. Pintore, M., Piras, G., Sotgiu, A., **Pintor**, **M.**, & Biggio, B. (2026). Evaluating line-level localization ability of learning-based code vulnerability detection models. Machine Learning, 115(4), 1–19. https://doi.org/10.1007/s10994-025-06902-1

3. Tony, C., **Pintor**, **M.**, Kretschmann, M., & Scandariato, R. (2026). Discrete prompt optimization using genetic algorithm for secure python code generation. Journal of Systems and Software, 232, 112682. https://doi.org/10.1016/j.jss.2025.112682

4. Angioni, D., Demetrio, L., **Pintor**, **M.**, Oneto, L., Anguita, D., Biggio, B., & Roli, F. (2025). Robustness-congruent adversarial training for secure machine learning model updates. IEEE Transactions on Pattern Analysis and Machine Intelligence, 47(9), 7457–7469. https://doi.org/10.1109/TPAMI.2025.3573237

5. Mura, R., Floris, G., Scionis, L., Piras, G., **Pintor**, **M.**, Demontis, A., Giacinto, G., Biggio, B., & Roli, F. (2025). Ho-fmn: Hyperparameter optimization for fast minimum-norm attacks. Neurocomputing, 616, 128918. https://doi.org/10.1016/j.neucom.2024.128918

6. Piras, G., **Pintor**, **M.**, Demontis, A., Biggio, B., Giacinto, G., & Roli, F. (2025). Adversarial pruning: A survey and benchmark of pruning methods for adversarial robustness. Pattern Recognition, 168, 1–13. https://doi.org/10.1016/j.patcog.2025.111788

7. Eghbal-zadeh, H., Zellinger, W., **Pintor**, **M.**, Grosse, K., Koutini, K., Moser, B. A., Biggio, B., & Widmer, G. (2024). Rethinking data augmentation for adversarial robustness. Information Sciences, 654, 119838. https://doi.org/10.1016/j.ins.2023.119838

8. Mirsky, Y., Demontis, A., Kotak, J., Shankar, R., Gelei, D., Yang, L., Zhang, X., Pintor, M., Lee, W., Elovici, Y., & Biggio, B. (2023). The threat of offensive ai to organizations. Computers & Security, 124, 1–23. https://doi.org/10.1016/j.cose.2022.103006

9. **Pintor**, **M.**, Angioni, D., Sotgiu, A., Demetrio, L., Demontis, A., Biggio, B., & Roli, F. (2023). Imagenet-patch: A dataset for benchmarking machine learning robustness against adversarial patches. Pattern Recognition, 134, 1–11. https://doi.org/10.1016/j.patcog.2022.109064

10. Zheng, Y., Feng, X., Xia, Z., Jiang, X., Demontis, A., **Pintor**, **M.**, Biggio, B., & Roli, F. (2023). Why adversarial reprogramming works, when it fails, and how to tell the difference. Information Sciences, 632, 130–143. https://doi.org/10.1016/j.ins.2023.02.086

11. Zheng, Y., Feng, X., Xia, Z., Jiang, X., **Pintor**, **M.**, Demontis, A., Biggio, B., & Roli, F. (2023). Stateful detection of adversarial reprogramming. Information Sciences, 642, 119093. https://doi.org/10.1016/j.ins.2023.119093

Top Conference Papers

1. Cinà, A. E., Rony, J., **Pintor**, **M.**, Demetrio, L., Demontis, A., Biggio, B., Ayed, I. B., & Roli, F. (2025). Attackbench: Evaluating gradient-based attacks for adversarial examples. AAAI-25 Technical Tracks 3, 39(3), 2600–2608. https://doi.org/10.1609/aaai.v39i3.32263

2. Cinà, A. E., Villani, F., **Pintor**, **M.**, Schönherr, L., Biggio, B., & Pelillo, M. (2025). σ-zero: Gradient-based optimization of ℓ_0-norm adversarial examples. International Conference on Representation Learning, 2025, 91199–91211.

3. **Pintor**, **M.**, Demetrio, L., Sotgiu, A., Demontis, A., Carlini, N., Biggio, B., & Roli, F. (2022). Indicators of attack failure: Debugging and improving optimization of adversarial examples. Advances in Neural Information Processing Systems 35 (NeurIPS 2022).

4. **Pintor**, **M.**, Roli, F., Brendel, W., & Biggio, B. (2021). Fast minimum-norm adversarial attacks through adaptive norm constraints. 35th Conference on Neural Information Processing Systems (NeurIPS 2021), 34, 20052–20062.

5. Demontis, A., Melis, M., **Pintor**, **M.**, Jagielski, M., Biggio, B., Oprea, A., Nita-Rotaru, C., & Roli, F. (2019). Why do adversarial attacks transfer? explaining transferability of evasion and poisoning attacks. 28th USENIX Security Symposium, USENIX Security 19, 321–338.

Preprints

1. Manca, C., Scano, C., Piras, G., Brau, F., **Pintor**, **M.**, & Biggio, B. (2026). Sage-5gc: Security-aware guidelines for evaluating anomaly detection in the 5g core network.

2. Minnei, L., Manca, C., Piras, G., Sotgiu, A., **Pintor**, **M.**, Ghiani, D., Maiorca, D., Giacinto, G., & Biggio, B. (2026). Label-efficient training updates for malware detection over time.

3. Piras, G., Mura, R., Brau, F., **Pintor**, **M.**, Oneto, L., Roli, F., & Biggio, B. (2026). Latent-space attacks for refusal evasion in language models.

4. Rodríguez-Lois, E., Brau, F., **Pintor**, **M.**, Biggio, B., & Pérez-González, F. (2026). Blackcatt: Black-box collusion aware traitor tracing in federated learning.

5. Eddoubi, H., Ricker, J., Cocchi, F., Baraldi, L., Sotgiu, A., **Pintor**, **M.**, Cornia, M., Baraldi, L., Fischer, A., Cucchiara, R., & Biggio, B. (2025). Raid: A dataset for testing the adversarial robustness of ai-generated image detectors.

6. Ghiani, D., Angioni, D., Piras, G., Sotgiu, A., Minnei, L., Gupta, S., **Pintor**, **M.**, Roli, F., & Biggio, B. (2025). Regression-aware continual learning for android malware detection.

7. Guo, W., Brau, F., **Pintor**, **M.**, Demontis, A., & Biggio, B. (2025). Silent until sparse: Backdoor attacks on semi-structured sparsity.

8. Guo, W., **Pintor**, **M.**, Demontis, A., & Biggio, B. (2025). Prototype-guided robust learning against backdoor attacks.

9. Gupta, S., Balia, R., Angioni, D., Brau, F., **Pintor**, **M.**, Demontis, A., Sebastian, A., Carta, S. M., Roli, F., & Biggio, B. (2025). Out-of-distribution detection for continual learning: Design principles and benchmarking.

10. Mura, R., Piras, G., Lukošiūtė, K., **Pintor**, **M.**, Karbasi, A., & Biggio, B. (2025). Latentbreak: Jailbreaking large language models through latent space feedback.

11. Pintore, M., **Pintor**, **M.**, Karatzas, D., & Biggio, B. (2025). Counterfeit answers: Adversarial forgery against ocr-free document visual question answering.

12. Piras, G., Zhao, Q., Brau, F., **Pintor**, **M.**, Wressnegger, C., & Biggio, B. (2025). S2ap: Score-space sharpness minimization for adversarial pruning.

13. Demontis, A., Gupta, S., **Pintor**, **M.**, Demetrio, L., Grosse, K., Lin, H.-Y., Fang, C., Biggio, B., & Roli, F. (2022). Security of deep reinforcement learning for autonomous driving: A survey.

14. **Pintor**, **M.**, Demetrio, L., Sotgiu, A., Melis, M., Demontis, A., & Biggio, B. (2019). Secml: A python library for secure and explainable machine learning.

Other Research Outputs

1. Towards debugging and improving adversarial robustness evaluations (Doctoral dissertation). (2022). Università degli Studi di Cagliari.

Last update: 2026/06/12