Lorenzo Cazzaro, Ph.D. student

lorenzo.cazzaro@unive.it @LorenzoCazz
https://lorenzocazzaro.github.io/
https://www.linkedin.com/in/lorenzo-cazzaro-622b37175/
https://github.com/LorenzoCazzaro
+393486494934
via J.Castelli 43, Venice, Italy

Education

09/2021 – · · · ·
- **Ph.D student in Computer Science, Ca' Foscari University of Venice**
  Research project title: *Principled Verification of Machine Learning Models*
  Supervisor: *prof. Stefano Calzavara*
  Research interests: *Adversarial Machine Learning, Verification of Machine Learning Models, Applications of Artificial Intelligence in Cybersecurity*.

03/2023 – 07/2023
- **Visiting Ph.D. student, CISPA Helmholtz Center for Information Security**
  Supervisor: *prof. Giancarlo Pellegrino*
  Research topic: *Improving Web Application Security through Artificial Intelligence.*

09/2022
- **Attended the CISPA Summer School on Trustworthy Artificial Intelligence, CISPA Helmholtz Center for Information Security**
  The CISPA Summer School on Trustworthy Artificial Intelligence covered different aspects of trustworthy Machine Learning like security, privacy and fairness.

11/2019 – 07/2021
- **M.Sc. in Computer Science - Software Dependability and Cyber Security (summa cum laude), Ca' Foscari University of Venice**
  Thesis title: *AMEBA: An Adaptive Approach to the Black-Box Evasion of Machine Learning Models.*

09/2016 – 11/2019
- **B.Sc. in Computer Science - Data Science (summa cum laude), Ca' Foscari University of Venice**
  Thesis title: *Transferability of Adversarial Examples from Linear SVM to Decision Tree Ensembles.*

Employment History

2022-2023
- **Teacher of the Laboratory on Advanced Artificial Intelligence: Linear Regression and Adversarial Machine Learning**, Ca' Foscari University of Venice.
- **Database Systems teaching assistant senior**, Ca' Foscari University of Venice.

2021-2023
- **Algorithms and Data Structures teaching assistant senior**, Ca' Foscari University of Venice.
- **Discrete Math teaching assistant**, Ca' Foscari University of Venice.

09/2020 - 01/2021
- **Linear Algebra teaching assistant**, Ca' Foscari University of Venice.

12/2019 - 03/2020
- **Research fellow in Adversarial Machine Learning**, Ca' Foscari University of Venice.

02/2019-03/2019
- **Trainee - Web Development**, Ennova Research S.r.l. - Mestre/Venice

Research Publications

Journal Papers

1. Calzavara, S., **Cazzaro, L.**, Lucchese, C., Marcuzzi, F., & Orlando, S. (2022). Beyond Robustness: Resilience Verification of Tree-Based Classifiers. *Computers & Security*, *121*, 102843.
   doi:https://doi.org/10.1016/j.cose.2022.102843

Conference Papers

1. Calzavara, S., **Cazzaro, L.**, Lucchese, C., & Marcuzzi, F. (2023). Explainable Global Fairness Verification of Tree-Based Classifiers. In *2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML 2023)* (pp. 1–17). doi:10.1109/SaTML54575.2023.00011
2. Calzavara, S., **Cazzaro, L.**, Pibiri, G. E., & Prezza, N. (2023). Verifiable Learning for Robust Tree Ensembles. In *ACM SIGSAC Conference on Computer and Communications Security (ACM CCS), Copenhagen, Denmark, November 26-30, 2023.* doi:10.48550/arXiv.2305.03626
3. Calzavara, S., **Cazzaro, L.**, & Lucchese, C. (2021). AMEBA: An Adaptive Approach to the Black-Box Evasion of Machine Learning Models. In J. Cao, M. H. Au, Z. Lin, & M. Yung (Eds.), *ASIA CCS '21: ACM Asia Conference on Computer and Communications Security, Virtual Event, Hong Kong, June 7-11, 2021* (pp. 292–306). doi:10.1145/3433210.3453114

Conference Presentations

2023 Speaker at **IEEE Conference on Secure and Trustworthy Machine Learning (IEEE SaTML 2023), Raleigh, North Carolina, USA** - Presentation of the paper *Explainable Global Fairness Verification of Tree-Based Classifiers.*

2022 Speaker at **AI for Security and Security of AI workshop (AISSAI22) in Italian Conference on Cybersecurity (ITASEC22), Rome, Italy** - Presentation of the short version of the paper *Beyond Robustness: Resilience Verification of Tree-Based Classifiers.*

2021 Speaker at **ACM Asia Conference on Computer and Communication Security (ASIACCS21), virtual event** - Presentation of the paper *AMEBA: An Adaptive Approach to the Black-Box Evasion of Machine Learning Models.*

Service

Artifact Evaluation committee I have been an Artifact Evaluator for the 17th Workshop On Offensive Technologies (WOOT '23) co-located with the 44th IEEE Symposium on Security and Privacy (IEEES&P 2023).

Invited reviewer I am or have been an anonymous reviewer for the following conferences/journals:

- **Conferences**: 31-th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN 2023) - Special session *Towards Machine Learning Models that We Can Trust: Testing, Improving, and Explaining Robustness*
- **Journals**: IEEE Transactions on Information Forensics and Security (IEEE TIFS) and Journal of Computer Security.

Service (continued)

External Reviewer

I performed some reviews as an external reviewer for:

- The 37-th Conference on Neural Information Processing Systems (NeurIPS 2023).
- The Italian Conference on Cybersecurity 2023 (ITASEC 2023).
- The 39-th Annual Computer Security Applications Conference (ACSAC 2022).
- The 6-th IEEE European Symposium on Security and Privacy (EuroS&P 2021).

Skills

Coding

Strong coding skills in C, C++ and Python; medium coding skills in LaTeX, R, Javascript, SQL

Web Dev

Experience with Angular, Apache Web Server, ExpressJS, Flask, PostgreSQL.

Machine Learning framework

Strong skills in using python for data cleaning and feature selection. Good knowledge of the packages scikit-learn and Tensorflow.

Research

Strong background in evasion attacks against Machine Learning models and robustness of Machine Learning algorithms. Good analytical and critical thinking and teamwork skills.

Projects

Fairness analyzer for decision tree ensembles

A fairness analyzer for decision tree ensembles written in C++. Given a decision tree ensemble and a set of sensitive features, it returns a set of logical formulas predicating on the subsets of instances on which it is guaranteed that the Machine Learning (ML) model does not perform unfairness (causal discrimination) on them. Link: `https://github.com/FedericoMarcuzzi/resilience-verification`.

Stability analyzer for decision tree ensembles

An analyzer for decision tree ensembles written in C++. Given a decision tree ensemble and an attack specification, it returns the regions of the feature space (hyper-rectangles) in which the ML model exhibits stability. Link: `https://github.com/LorenzoCazzaro/explainable-global-fairness-verification`.