Ioan-Tudor Cebere

Email: tudorcebere@gmail.com **Personal Webpage**

Research interests

Differential Privacy, Machine Learning, Privacy-Preserving Machine Learning, Membership Inference, Reconstruction Attacks

Education

Inria Montpellier, France
PhD in Computer Science (expected graduation 12.2026) 11.2022 – Present
Advisor: Aurélien Bellet.

Ecole Normale Supérieure de Lyon Lyon, France
MSc in Theoretical Computer Science 09.2021 – 07.2022
Advisor: Sonia Ben Mokhtar
Thesis: *Blackbox Membership Inference Attack via Memorization.*

Politehnica University of Bucharest Bucharest, Romania
BSc in Computer Science and Engineering 10.2016 – 07.2020
Supervisor: Matei Popovici
Thesis: *ROS simulator for Reinforcement Learning*

Publications

Privacy in Theory, Bugs in Practice: Grey-Box Auditing of Differential Privacy Libraries
TC, David Erb, Damien Desfontaines, Aurélien Bellet, Jack Fitzsimons
PoPETS 2026.3

Membership Inference Attacks from Causal Principles
Mathieu Even, Clément Berenfeld, Linus Bleistein, TC, Julie Josse, Aurélien Bellet
Under Review for ICML 2026

Rate-Constrained Optimization with Differential Privacy
Mohammad Yaghini*, TC*, Mike Menart, Aurélien Bellet, Nicolas Papernot
ICLR 2026, equal contribution

Tighter Privacy Auditing of DP-SGD in the Hidden State Threat Model
TC, Aurélien Bellet, Nicolas Papernot
ICLR 2025

Confidential-DPproof: Confidential Proof of DP Training
AS Shamsabadi, Gefei Tan, TC, Aurélien Bellet, et al.
ICLR 2024 ***spotlight paper***

Workshop Papers

Syft 0.5: A platform for universally deployable structured transparency
Adam James Hall, Madhava Jay, TC, Bogdan Cebere et al.
Distributed and Private Machine Learning (DPML), ICLR Workshop, 2021.

PyVertical: A Vertical FL Framework for Multi-headed SplitNN
Daniele Romanini, Adam J. Hall, ..., TC et al.
Distributed and Private Machine Learning (DPML), ICLR Workshop, 2021.

Awards and scholarships

Google Fellowship in Security and Privacy	2025
Ampère Scholarship of Excellence (12.000 €)	2021
1^{st} place PyTorch Hackaton, Facebook (5.000 $)	2020
1^{st} place, Machine Learning Contest, Cognizant (25.000 $)	2019

Work experience

OpenDP, Harvard University Cambridge, Massachusetts
Associate *9.2024 - present*
Review code and proofs in the OpenDP library. Co-lead a workgroup on privacy attacks, helping practitioners understand data privacy threat models.
Reference: Michael Shoemate

Intern *6.2024 - 9.2024*
Implemented support for the Bounded Range differential privacy variant, enhancing privacy guarantees in core OpenDP components like noisy argmax and the exponential mechanism. Reviewed and proofread implementations of various building blocks.
Reference: Michael Shoemate

Vector Institute, University of Toronto Toronto, Canada
Research Intern *10.2023 - 02.2024*
Investigated privacy guarantees of differentially private SGD (DP-SGD) in a relaxed threat model called the hidden state. We were able to show an important distinction between convex and non-convex models in this threat model.
Reference: Nicolas Papernot

OpenMined, Syft Library Remote
Core Engineer *02.2020 - 12.2022*
Syft is a library that aims to make machine learning privacy-friendly. Enhanced the performance and security of the distributed learning stack. Designed a JAX tensor type to track information necessary for individual differential privacy.
Reference: Andrew Trask

Inria Lyon, France
MSc internship *02.2022 - 07.2022*
Developed a novel membership inference attack exploiting model misclassifications, eliminating the need for costly shadow models. The attack is time-efficient and suitable for privacy hypothesis testing.
Reference: Sonia Ben Mokhtar

Mines Paristech, PSL University Sophia Antipolis, France
Research Engineering intern *06.2020 - 01.2021*
Developed a task scheduler for a fleet of robots using graph modeling and reinforcement learning, ensuring operation under various fault models.
Reference: Sebastien Travadel

UiPath Bucharest, Romania
Machine Learning Engineering intern *06.2019 - 09.2019*
Developed a recommendation system using a MultiVAE architecture for collaborative filtering. Engineered a document denoising tool utilizing CycleGANs to enhance OCR accuracy.

Invited Talks

Rate-Constrained Private Optimization
Inria Lille, October 2025;
Tighter privacy auditing in the Hidden State
Microsoft Research, December 2023; Google Deepmind, August 2024; Inria Montpellier, November 2024; National University of Singapore, May 2025;
Privacy Auditing & Privacy Amplifications
Inria Lille, April 2023; University of Toronto, November 2023;
Syft 0.5: A platform for universally deployable structured transparency
Distributed and Private Machine Learning (DPML), ICLR Workshop, 2021

Student Mentorship

David ERB (TUM) - Research Internship & BSc Thesis	Fall '25 - Summer '26
Vaibhav Vardhan (IIT Kanpur) - Google Summer of Code	Summer 2021
Param Mirani (VIT Pune) - Google Summer of Code	Summer 2021
Aditi Verma (IITS) - Google Summer of Code	Summer 2021

Service and outreach

Reviewer
ICLR 2026, SatML 2026, TPDP 2025, ICML 2025, AISTATS 2025 (emergency), ICLR 2025, ICML 2024, NeurIPS 2023, ICML 2023 (reduced load)

Romanian Open Source Educational 03.2018 – present
Led projects supporting the Romanian open-source community, impacting over 2,000 students through courses, talks, and workshops. Personally trained over 200 students. Reference: Razvan Deaconescu

OpenMined Research 12.2022 – 12.2023

Organized monthly meetings, seminars, and reading groups on privacy, security, robustness, and fairness in machine learning, fostering community engagement in trustworthy AI.

Reference: Andrew Trask