

ADAM DZIEDZIC

PERSONAL DETAILS

PHONE: +1 872 222 8183
EMAIL: adam.dziedzic@sprintml.com
PERSONAL WEB PAGE: <https://adam-dziedzic.com/>
LINKEDIN: <https://www.linkedin.com/in/adziedzic>
GITHUB: <https://github.com/adam-dziedzic>
TWITTER: https://twitter.com/adam_dziedzic

ACADEMIC AND RESEARCH EXPERIENCE

CURRENT **CISPA Helmholtz Center for Information Security**, Germany
SEPTEMBER 2023 *Tenure Track Faculty Member*
My research is focused on secure and trustworthy Machine Learning as a Service (MLaaS). I design robust and reliable machine learning methods for training and inference of ML models while preserving data privacy and model confidentiality.
GROUP: [SprintML Lab](#)

SEPTEMBER 2023 **Vector Institute & the University of Toronto**, Canada
SEPTEMBER 2020 *Postdoctoral Fellow in COMPUTER SCIENCE*
GROUP: [CleverHans Lab](#)
ADVISOR: [Professor Nicolas PAPERNOT](#)
RESEARCH AREAS: Trustworthy & Collaborative Machine Learning

SEPTEMBER 2017 **Google** (MADISON, USA)
JUNE 2017 *PhD Software Engineering Intern at Data Infrastructure and Analysis Team*
Mentor: Goetz Graefe

JUNE 2017 **Microsoft Research** (REDMOND, USA)
MARCH 2017 *Research Intern at Data Management, Exploration and Mining (DMX)*
Mentors: Vivek Narasayya and Sudipto Das.

JUNE 2015 **École Polytechnique Fédérale de Lausanne (EPFL)**, Switzerland
OCTOBER 2014 *Research Intern at Data Intensive Applications and Systems*
ADVISOR: Professor Anastasia AILAMAKI

DECEMBER 2012 **CERN** (GENEVA, SWITZERLAND)
APRIL 2012 *Technical Student at IT Department and CERN Computer Center*

EDUCATION

AUGUST 2020 **University of Chicago**, USA
JULY 2015 PhD Program in COMPUTER SCIENCE
ADVISOR: [Professor Sanjay KRISHNAN](#)
RESEARCH AREAS: Robust Machine & Deep Learning, Data Analysis and Database Systems
GPA: 3.96/4

TEACHING ASSISTANT: [Fundamentals of Deep Learning](#), Introduction to Databases, Databases for Public Policy

- SEPTEMBER 2014 **Warsaw University of Technology**, Poland
OCTOBER 2013 Graduate Research Assistant
MAJOR: Computer Information System Engineering
MAIN TOPIC: Big Data
ADVISOR: Professor Jan MULAWKA
TEACHING ASSISTANT: Bioinformatics Algorithms
- MARCH 2013 **Warsaw University of Technology**, Poland
OCTOBER 2011 Master of Science in COMPUTER SCIENCE
MAJOR: Computer Information System Engineering
THESIS: “An analysis and comparison of non-relational (NoSQL) databases with an example of application using CouchDB.”
ADVISOR: Professor Piotr GAWRYSIAK
GPA: 4.93/5 (top 5%) THE FINAL GRADE: Excellent
- SEPTEMBER 2011 **Warsaw University of Technology**, Poland
FEBRUARY 2011 Bachelor of Science in COMPUTER SCIENCE
MAJOR: Computer Information System Engineering
THESIS: “Document management system – application in three-tiered architecture.”
ADVISOR: Ph.D. Eng Jarosław DAWIDCZYK
GPA: 4.80/5 (top 5%) THE FINAL GRADE: Excellent
- JANUARY 2011 **Technical University of Denmark**
AUGUST 2010 GPA: 11.71/12
- JUNE 2010 **Warsaw University of Technology**, Poland
OCTOBER 2007 MAJOR: Computer Information System Engineering

WORK EXPERIENCE

- | | |
|----------------|---|
| AUGUST 2013 | Barclays Investment Bank (LONDON, THE UK) |
| JUNE 2013 | <i>Analyst at Equities Derivatives Technology</i> |
| MARCH 2012 | Mobile Startup |
| APRIL 2012 | Application providing aspects of music social interactions |
| JULY 2010 | Tekten SP. Z O.O. (WARSAW, POLAND)
<i>Database designer, Java and PL/SQL software developer</i>
Telecom System Project |
| SEPTEMBER 2009 | Torn SP. Z O.O. (WARSAW, POLAND) |
| JULY 2009 | <i>Java and JavaScript software developer</i>
Financial and accounting system project |

AWARDS

- 2023 The Best Poster Award at the MLinPL (Machine Learning in Poland) Conference for our work on *Bucks for Buckets (B4B): Active Defenses Against Stealing Encoders*.
- 2022 Highlighted Paper on a new defense against model extraction at International Conference on Learning Representations (ICLR).
- 2022 Highlighted Reviewer at the International Conference on Learning Representations (ICLR).
- 2019 Travel Award at International Conference on Machine Learning (ICML).
- 2018 Travel Award at SIGMOD (Special Interest Group on Management of Data).
- 2011-2012 The scholarship of the Rector of the Warsaw University of Technology for my achievements during the Master’s program.
- 2007-2011 The academic scholarship for the best faculty students (granted on a yearly basis and based on GPA).

PUBLICATIONS

- eBioMedicine 2024 Congyu Fang, Adam Dziedzic, Lin Zhang, Laura Oliva, Amol Verma, Fahad Razak, Nicolas Papernot, Bo Wang *Decentralised, Collaborative, and Privacy-preserving Machine Learning for Multi-Hospital Data*
- ICLR 2024 Wenhao Wang, Muhammad Ahmad Kaleem, Adam Dziedzic, Michael Backes, Nicolas Papernot, Franziska Boenisch *Memorization in Self-Supervised Learning Improves Downstream Generalization*
- NeurIPS 2023 Jan Dubiński, Stanisław Pawlak, Franziska Boenisch, Tomasz Trzcinski, Adam Dziedzic *Bucks for Buckets (B4B): Active Defenses Against Stealing Encoders*
- NeurIPS 2023 Nicholas Franzese, Adam Dziedzic, Christopher A. Choquette-Choo, Mark R. Thomas, Muhammad Ahmad Kaleem, Stephan Rabanser, Congyu Fang, Somesh Jha, Nicolas Papernot, Xiao Wang *Robust and Actively Secure Serverless Collaborative Learning*
- NeurIPS 2023 Haonan Duan, Adam Dziedzic, Nicolas Papernot, Franziska Boenisch *Flocks of Stochastic Parrots: Differentially Private Prompt Learning for Large Language Models*
- NeurIPS 2023 Franziska Boenisch, Christopher Mühl, Adam Dziedzic, Roy Rinberg, Nicolas Papernot *Have it your way: Individualized Privacy Assignment for DP-SGD*
- PETs 2023 Adam Dziedzic, Christopher A. Choquette-Choo, Natalie Dullerud, Vinith Suriyakumar, Ali Shahin Shamsabadi, Muhammad Ahmad Kaleem, Somesh Jha, Nicolas Papernot, Xiao Wang *Private Multi-Winner Voting for Machine Learning*
- PETs 2023 Franziska Boenisch, Christopher Mühl, Roy Rinberg, Jannis Ihrig, Adam Dziedzic *Individualized PATE: Differentially Private Machine Learning with Individual Privacy Guarantees*

EuroS&P 2023	Franziska Boenisch, Adam Dziedzic, Roei Schuster, Ali Shahin Shamsabadi, Iliia Shumailov, Nicolas Papernot <i>When the Curious Abandon Honesty: Federated Learning Is Not Private</i>
EuroS&P 2023	Franziska Boenisch, Adam Dziedzic, Roei Schuster, Ali Shahin Shamsabadi, Iliia Shumailov, Nicolas Papernot <i>Is Federated Learning a Practical PET Yet?</i>
ICLR Workshop 2023	Adam Dziedzic, Franziska Boenisch, Mingjian Jiang, Haonan Duan, Nicolas Papernot <i>Sentence Embedding Encoders are Easy to Steal but Hard to Defend</i> (ICLR 2023 Workshop on Trustworthy ML)
ICLR 2022	Adam Dziedzic, Muhammad Ahmad Kaleem, Yu Shen Lu, Nicolas Papernot <i>Increasing the Cost of Model Extraction with Calibrated Proof of Work</i> SPOTLIGHT (top 5% of accepted papers)
NeurIPS 2022	Adam Dziedzic, Haonan Duan, Muhammad Ahmad Kaleem, Nikita Dhawan, Jonas Guan, Yannis Cattan, Franziska Boenisch, Nicolas Papernot <i>Dataset Inference for Self-Supervised Models</i>
ICML 2022	Adam Dziedzic, Nikita Dhawan, Muhammad Ahmad Kaleem, Jonas Guan, Nicolas Papernot <i>On the Difficulty of Defending Self-Supervised Learning against Model Extraction</i>
ICLR 2021	Christopher A. Choquette-Choo, Natalie Dullerud, Adam Dziedzic, Yunxiang Zhang, Somesh Jha, Nicolas Papernot, Xiao Wang <i>CaPC Learning: Confidential and Private Collaborative Learning</i>
ACL 2020	Dan Hendrycks, Xiaoyuan Liu, Eric Wallace, Adam Dziedzic, Rishabh Krishnan, Dawn Song <i>Pretrained Transformers Improve Out-of-Distribution Robustness</i>
JOR 2020	Arnold Wong, Garrett Harada, Remy Lee, Sapan D. Gandhi, Adam Dziedzic, Alejandro Espinoza-Orias, Mohamad Parnianpour, Philip Louie, Bryce Basques, Howard S. An, Dino Samartzis <i>Preoperative paraspinal neck muscle characteristics predict early-onset adjacent segment degeneration in anterior cervical fusion patients: a machine-learning modeling analysis</i>
OJVT 2020	Adam Dziedzic, Vanlin Sathya, Monisha Ghosh, Sanjay Krishnan <i>Machine Learning for Fair Spectrum Sharing in Dense LTE Wi-Fi Coexistence</i>
ICNC 2020	Vanlin Sathya, Adam Dziedzic, Monisha Ghosh, Sanjay Krishnan <i>Machine learning-based detection of multiple Wi-Fi BSSs for LTE-U CSAT</i>
Ph.D. 2020	Adam Dziedzic <i>Input and Model Compression for Adaptive and Robust Neural Networks</i> (Ph.D. Thesis)
ICML 2019	Adam Dziedzic, Ioannis Paparrizos, Sanjay Krishnan, Aaron J. Elmore, Michael Franklin <i>Band-limited Training and Inference for Convolutional Neural Networks</i> (paper) code: https://github.com/adam-dziedzic/bandlimited-cnns

SIGOPS 2019 Sanjay Krishnan, Aaron J. Elmore, Michael Franklin, Ioannis Paparrizos, Zechao Shang, Adam Dzielicz, Rui Liu *Artificial Intelligence in Resource-Constrained and Shared Environments*

CIDR 2019 Sanjay Krishnan, Adam Dzielicz, Aaron J. Elmore *DeepLens: Towards a Visual Data Management System*

SIGMOD 2018 Adam Dzielicz, Jingjing Wang, Sudipto Das, Bolin Ding, Vivek R. Narasayya, Manoj Syamala *Columnstore and B+ tree – Are Hybrid Physical Designs Important?*

UChicago 2017 Adam Dzielicz *Data Loading, Transformation, and Migration for Database Management Systems* (Master’s thesis)

CIDR 2017 Tim Mattson, Vijay Gadepally, Zuohao She, Adam Dzielicz, Jeff Parkhurst *Demonstrating the BigDAWG Polystore System for Ocean Metagenomic Analysis*

VLDB ADMS 2016 Adam Dzielicz, Manos Karpathiotakis, Ioannis Alagiannis, Raja Appuswamy, Anastasia Ailamaki *DBMS Data Loading: An Analysis on Modern Hardware*

HPEC 2016 Adam Dzielicz, Aaron J. Elmore, Michael Stonebraker *Data Transformation and Migration in Polystores* (paper) code: <https://github.com/bigdawg-istc/bigdawg>

HPEC 2016 John Meehan, Stan Zdonik, Shaobo Tian, Yulong Tian, Nesime Tatbul, Adam Dzielicz and Aaron J. Elmore *Integrating Real-Time and Batch Processing in a Polystore*

IEEE VIS DSIA 2015 Adam Dzielicz, Jennie Duggan, Aaron J. Elmore, Vijay Gadepally, Michael Stonebraker *BigDAWG: a Polystore for Diverse Interactive Applications*

SPIE 2014 Adam Dzielicz, Jan Mulawka. *Analysis and Comparison of databases with an introduction to consistent references in big data storage systems*

PREPRINTS

ArXiv 2022 Stephan Rabanser, Anvith Thudi, Kimia Hamidieh, Adam Dzielicz, Nicolas Papernot *Selective Classification Via Neural Network Training Dynamics*

ArXiv 2022 Adam Dzielicz, Stephan Rabanser, Mohammad Yaghini, Armin Ale, Murat A. Erdogdu, Nicolas Papernot *p-DkNN: Out-of-Distribution Detection Through Statistical Testing of Deep Representations*

ArXiv 2021 Adelin Travers, Lorna Licollari, Guanghan Wang, Varun Chandrasekaran, Adam Dzielicz, David Lie, Nicolas Papernot *On the Exploitability of Audio Machine Learning Pipelines to Surreptitious Adversarial Examples*

Intel 2021	Ahmad-Reza Sadeghi, Ferdinand Brasser, Markus Miettinen, Thien Duc Nguyen, Thomas Given-Wilson, Axel Legay, Murali Annaaram, Salman Avestimeh, Alexandra Dmitrienko, Farinaz Koushanfar, Buse Gul Atli, Florian Kerschbaum, Lachlan J. Gunn, N. Asokan, Matthias Schunter, Rosario Cammarota, Adam Dziedzic, Nicolas Papernot, Virginia Smith, Reza Shokri <i>Private AI Collaborative Research Institute: Vision, Challenges, and Opportunities</i>
ArXiv 2020	Adam Dziedzic, Sanjay Krishnan <i>Empirical Evaluation of Perturbation-based Defenses</i>

TEACHING

Trustworthy Machine Learning 2024	<i>TML</i> : Advanced Lecture on Trustworthy Machine Learning at the Saarland University and CISPA given at the summer semester 2024. This course explores the different aspects of trustworthy machine learning, including Privacy, Collaborative Learning, Model Confidentiality, Robustness, Fairness and Bias, Explainability, Security, and Governance.
Trustworthy Machine Learning 2023	<i>TrustML</i> : Seminar on Trustworthy Machine Learning at the Saarland University and CISPA given in the winter semester 2023/2024. The main focus of the seminar is on the security, privacy, confidentiality, and robustness of machine learning models.
Deep Learning	<i>TTIC-31230</i> : Teaching assistant for the course on <i>Fundamentals of Deep Learning</i> taught by <i>Prof. David McAllester</i> (Winter 2020)
Database Systems	<i>CS23500/33550</i> : Teaching assistant for the course on <i>Database Systems</i> taught by <i>Prof. Aaron J. Elmore</i> (Autumn 2015, Spring 2016, Winter 2017, Winter 2018, Spring 2018)
Bioinformatics Algorithms	<i>MBI</i> : Teaching assistant for the course on <i>Methods in Bioinformatics</i> taught by <i>Prof. Robert M. Nowak</i> (Spring 2014)

TALKS

2024	Gave a seminar on <i>Private Prompt Learning for Large Language Models</i> at the University of Waterloo (April 12).
2023	Talk about our paper on <i>Flocks of Stochastic Parrots: Differentially Private Prompt Learning for Large Language Models</i> at IDEAS NCBR (December).
2023	Talk on model stealing and defenses and speaker at the panel discussion at the NeurIPS workshop on <i>Backdoors in Deep Learning: The Good, the Bad, and the Ugly</i> .
2023	Presentation of our paper on <i>Flocks of Stochastic Parrots: Differentially Private Prompt Learning for Large Language Models</i> at MLinPL 2023 (October 27th).
2023	Presentation of our paper on <i>Private Multi-Winner Voting for Machine Learning</i> at PETS 2023 (July 12th).
2023	Talk on "Is this model mine? On stealing and defending machine learning models." for ML capstone class ECE697 from the University of Wisconsin-Madison (June 20th).
2023	Podcast Interview: On model stealing and defenses hosted by Matt Faltyn. https://traincheck.buzzsprout.com/

- 2023 Is this model mine? On stealing and defending machine learning models. CISPA, Germany.
- 2022 Is this Encoder Mine? On Stealing and Defending Self-Supervised Encoders
AI Safety Unconference NeurIPS 2022
- 2022 Stealing and Defending Self-Supervised Models. Invited talk on Dataset Inference for Self-Supervised Models at **ML Collective** (a nonprofit research organization) during their reading group "Deep Learning: Classics and Trends" which runs weekly, is fully virtual and is open to the public. They have 3000+ email subscribers and on average 100 weekly attendees.
- 2022 **Managing AI Risk - Cybersecurity & Data Risk Workstream at Vector Institute:** Is this Encoder Mine? On Stealing and Defending Self-Supervised Encoders
- 2022 Is this model mine? On stealing and defending machine learning models
University of Michigan at Ann Arbor
- 2022 Collaborative Machine Learning.
Vector Talk Series
- 2021 Confidential and Private Collaborative Learning.
Scotia Bank - Research Frontier Talk Series
- 2021 CaPC Learning: Confidential and Private Collaborative Learning.
Vector School: AI Model Governance
- 2021 CaPC Learning: Confidential and Private Collaborative Learning.
Invited Speaker for the Third Workshop on Privacy in Natural Language Processing.
- 2021 CaPC Learning: Confidential and Private Collaborative Learning.
The MLFL series, hosted by the Center for Data Science, UMass Amherst.
- 2021 CaPC Learning: Confidential and Private Collaborative Learning.
Flow Seminar
- 2021 CaPC Learning: Confidential and Private Collaborative Learning.
Intel Labs
- 2020 CaPC Learning: Confidential and Private Collaborative Learning.
Vector Institute
- 2018 Columnstore and B+ tree – are hybrid physical designs important?
University of California, Berkeley
- 2018 Columnstore and B+ tree – are hybrid physical designs important?
Imperial College London
- 2018 Columnstore and B+ tree – are hybrid physical designs important?
Oracle
- 2018 Columnstore and B+ tree – are hybrid physical designs important?
Microsoft Research
- 2018 Columnstore and B+ tree – are hybrid physical designs important?
MemSQL

SERVICE AND VOLUNTEERING

- Hackathon March 2024 Co-Organized [a hackathon in Warsaw, Poland in March 2024](#). There were 120 participants and the winning team was offered internships at SprintML Lab. The tasks at the hackathon were based on our research on [model stealing and defenses](#).
- SaTML Served as a session chair at the SaTML 2024 conference for the [session on Collaborative learning](#). Reviewer at the conference: 2023, 2024.

Vector	Served on the Research Adjudication Committee for the Vector Scholarship in Artificial Intelligence : 2022.
USENIX	Program Committee Member: 2023.
CCS	Program Committee Member: 2023 in the Machine Learning and Security Track , 2024 in the Machine Learning and Security Track .
ICLR	Reviewer at the International Conference on Learning Representations: 2019, 2020, 2021, 2022 highlighted reviewer, top 5% , 2023, 2024.
ICML	Reviewer at the International Conference on Machine Learning: 2021, 2022, 2023, 2024.
NeurIPS	Reviewer at the conference on Neural Information Processing Systems: 2021, 2022, 2023, 2024.

REFERENCES

NICOLAS PAPERNOT	Assistant Professor at the University of Toronto and the Vector Institute EMAIL: nicolas.papernot@utoronto.ca
SANJAY KRISHNAN	Assistant Professor at the University of Chicago EMAIL: skr@uchicago.edu
SOMESH JHA	Lubar Professor at the University of Wisconsin, Madison EMAIL: jha@cs.wisc.edu
XIAO WANG	Assistant Professor at Northwestern University EMAIL: wangxiao@cs.northwestern.edu
VIVEK NARASAYYA	Principal Researcher at Microsoft Research, Redmond EMAIL: viveknar@microsoft.com
MICHAEL FRANKLIN	Liew Family Chairman of Computer Science at the University of Chicago EMAIL: mjfranklin@uchicago.edu