

Kexin Li (Cassie)

Curriculum Vitae

Toronto, Canada — cassiekx.li at amail.utoronto.ca — cassiekexinli.com — GitHub — LinkedIn

EDUCATION & EMPLOYMENT HISTORY

Doctor of Philosophy in Computer Engineering

September 2024 – Present

Department of Electrical & Computer Engineering, **University of Toronto**

Research Area: Trustworthy Machine Learning and Systems Security, with a focus on building secure and privacy-preserving Machine Learning systems, watermarking for intellectual property protection and detection of AI-generated content, and adversarial robustness.

Supervisor: Prof. David Lie

Grade: A+

Master of Applied Science in Computer Engineering

September 2022 – August 2024

Department of Electrical & Computer Engineering, **University of Toronto**

Thesis: *Recovering Utility in LDP Schemes by Training with Noise²*: Invented methods to improve model utility under local differential privacy constraints while preserving strong privacy guarantees.

Supervisor: Prof. David Lie

Grade: A+

Bachelor of Applied Science with High Honours

September 2017 – June 2022

Department of Electrical & Computer Engineering, **University of Toronto**

Computer Engineering Specialist; Minor in Artificial Intelligence

CGPA: 3.91 / 4.0

Software Engineer

May 2020 – May 2021

Intel Corporation (now Altera), Toronto, Canada

- Enhanced production LLVM-based High-Level Synthesis (HLS) compiler toolchains (oneAPI, OpenCL, HLS compiler) for Intel FPGAs, improving performance and robustness for internal and external users.
- Shipped compiler features and analysis passes that supported scalable compilation of complex FPGA workloads across multiple device families.
- Executed weekly Quality-of-Results (QoR) analysis, diagnosing performance and resource regressions across compiler versions and architectures to ensure production readiness.
- Collaborated with globally distributed teams to assess architectural changes and clearly communicate performance, reliability, and usability trade-offs to stakeholders and customers.

PUBLICATIONS

Conference and Workshop Articles

- Guozhen Ding, **Kexin Li**, Ilya Grishchenko, and David Lie, *RFC2TLA+: Extracting and Verifying Formal Models from RFC Documents using Continuous LLM Feedback*, in Proceedings of the 41st IEEE/ACM International Conference on Automated Software Engineering (ASE 2026), Munich, Germany, Oct. 2026 (*To Appear*). (**PhD Work**)
- **Kexin Li**, Aastha Mehta, and David Lie *LDPKiT: Superimposing Remote Queries for Privacy-Preserving Local Model Training*, in Proceedings of 6th International Workshop on Advances on Security and Privacy Technologies and Solutions (IWAPS 2026), in conjunction with ARES 2026, Linköping, Sweden, Aug. 2026, Springer LNCS (*To Appear*). (**Master's Work**)
- Yunpeng Liu, **Kexin Li**, Zhuotao Liu, Bihan Wen, Ke Xu, Weiqiang Wang, Wenbiao Zhao, and Qi Li. 2023. *Provenance of Training without Training Data: Towards Privacy-Preserving DNN Model Ownership Verification*. In Proceedings of the ACM Web Conference 2023 (WWW '23). Association for Computing Machinery, New York, NY, USA, 1980–1990. <https://doi.org/10.1145/3543507.3583198>
- Yu Ting Chen, Jin Hee Kim, **Kexin Li**, Graham Hoyes, Jason H. Anderson. *High-Level Synthesis Techniques to Generate Deeply Pipelined Circuits for FPGAs with Registered Routing*. IEEE International Conference on Field-Programmable Technology (ICFPT), 2019.

Preprints

- **Kexin Li***, Xiao Hu*, Ilya Grishchenko, and David Lie, *LambdaMark: Semantic Audio Watermarking for Robustness and Radioactivity*, 2026 (**PhD Work**)
- Guozhen Ding, **Kexin Li**, Ilya Grishchenko, and David Lie *RFC2TLA+: Extracting and Verifying Formal Models from RFC Documents using Continuous LLM Feedback*, 2026. (**PhD Work**)
- **Kexin Li**, Guozhen Ding, Ilya Grishchenko, and David Lie. *HMARK: Radioactive Multi-Bit Semantic-Latent Watermarking for Diffusion Models*. arXiv, 2025. (**PhD Work**)
- **Kexin Li***, Xiao Hu*, Ilya Grishchenko, and David Lie. *HarmonicAttack: An Adaptive Cross-Domain Audio Watermark Removal*. arXiv, 2025. (**PhD Work**)

OTHER RESEARCH EXPERIENCE

Research Intern

May 2021 – April 2022

Institute for Network Sciences and Cyberspace, **Tsinghua University**

Supervisor: Prof. Qi Li

- Designed a novel intellectual property protection mechanism for deep neural networks (DNNs), targeting robustness against model extraction and unauthorized reuse.
- Characterized attacker behaviors and identified security limitations in existing DNN protection techniques through empirical analysis.
- Conducted systematic, large-scale empirical evaluations across diverse adversarial threat models to assess defense effectiveness.

Summer Research Intern

May 2019 – August 2019

Programmable Digital Systems Group, **University of Toronto**

Supervisor: Prof. Jason H. Anderson

- Enhanced the LegUp High-Level Synthesis framework to exploit register-rich FPGA architectures, enabling deeper pipelines and higher performance designs.
- Implemented LLVM backend extensions, improving the pipeline depth and throughput of synthesized Intel Stratix 10 FPGA circuits.

GRANTS, FELLOWSHIPS & AWARDS

- Queen Elizabeth II Graduate Scholarship in Science & Technology (2026 - Present)
- SRI Graduate Fellowship, Schwartz Reisman Institute (2025 – 2026)
- University of Toronto Fellowship (2022 – Present)
- Certificate of Distinction, Capstone Project, University of Toronto (2022)
- Dean's Honour List, University of Toronto (2017 – 2022)
- University of Toronto Excellence Summer Research Award (2019) [4 recipients]

PROFESSIONAL SERVICE & AFFILIATIONS

- Program Committee Member, AI4Good Workshop at ICML 2026 and NeurIPS 2026
- Reviewer, AI4Good Workshop at ICML 2026 and NeurIPS 2026
- Graduate Fellow Affiliate, Schwartz Reisman Institute (2026 – Present)
- Faculty Affiliate Researcher, Vector Institute (2023 – Present)
- Graduate Fellow, Schwartz Reisman Institute (2025 – 2026)
- NeurIPS Conference Ethics Reviewer (2024, 2025, 2026)
- NeurIPS Datasets and Benchmarks Track Ethics Reviewer (2025, 2026)

TEACHING & ADVISING

Teaching Assistantships

- **ECE1508H1 Deep Generative Models**, University of Toronto (2026 – Present)
- **ECE1508H1 Applied Deep Learning**, University of Toronto (2025 – Present)
- **ECE568H1 Computer Security**, University of Toronto (2023 – Present)
- **ECE244H1 Programming Fundamentals**, University of Toronto (2022 – Present)

TECHNICAL SKILLS

- **Research Areas:** Trustworthy Machine Learning, Privacy-Preserving Machine Learning, Watermarking, Systems Security
- **Programming Languages:** C, C++, Python, Java, ARM Assembly, JavaScript, SQL
- **Systems & Tools:** LLVM, Linux, Intel oneAPI, Verilog, Quartus, ModelSim, MATLAB
- **Languages:** English, Mandarin, Korean, French