ALEX EDMONDS

aedmonds@cs.toronto.edu
www.cs.toronto.edu/~edmonds

CURRICULUM VITAE

ABOUT

I completed my PhD at the University of Toronto, supervised by Aleksandar Nikolov and Toni Pitassi. The focus of my research was on differential privacy, particularly the local model, with connections to learning theory, information theory, and duality. My research with my collaborators gave sample-complexity characterizations under local differential privacy of fundamental tasks such as linear query release, agnostic PAC learning, and realizable PAC learning.

Currently, I am interested in moving into the applied side of data privacy. I can bring a strong theoretical and analytical grounding to the task of putting privacy techiques into practice.

INTERESTS

Differential privacy, learning theory, information theory.

EDUCATION

University of Toronto 2017 - 2023
Doctor of Philosophy – *Computer Science*

Simons Institute for the Theory of Computing 2019
Visiting Graduate Student – *Data Privacy: Foundations and Applications*

Simons Institute for the Theory of Computing 2017
Visiting Graduate Student – *Foundations of Machine Learning*

University of Toronto 2015 - 2017
Master of Science – *Computer Science*

University of Toronto 2008 - 2014
Honours Bachelor of Science – *Mathematics Specialist Program*

ACADEMIC AWARDS

PhD Thesis recommended for departmental award by external examiner
University College Galois Mathematics Scholarship 2014
University College Alumni Association Scholarship 2012
George Roderick Fraser Scholarship for Mathematical Studies 2008

RESEARCH

Edmonds, A. (2023). "Sample-Complexity Optimality Under Local Differential Privacy and Related Models (PhD Thesis)".
http://www.cs.toronto.edu/ edmonds/doc/phd-thesis.pdf.

My thesis gives a unified presentation of the results from [ENU19] and [ENP22]

which consider the problems of characterizing the sample-complexity under non-interactive local privacy of linear query release as well as both agnostic and realizable variants of PAC learning and refutation. In addition, my thesis presents joint research with Nikolov and Pitassi which derives a polynomial equivalence between sample complexities for learning under single-intrusion pan-privacy and learning under sequentially interactive local privacy. My thesis also presents research with Nikolov and Pitassi which derives a polynomial equivalence between the query complexity of learning under "correlational" linear queries and sample complexity of learning under non-interactive local privacy.

Edmonds, A., A. Nikolov, T. Pitassi (2022). "Learning versus Refutation in Noninteractive Local Differential Privacy". NeurIPS 2022.
https://arxiv.org/abs/2210.15439.

We consider two basic tasks to be performed under non-interactive local differential privacy (LDP): learning and refutation. Whereas our previous work showed that agnostic refutation implies agnostic learnability under non-interactive LDP, it left open the question of the converse. This work resolves this question, showing these problems are equivalent by way of a complete characterization of the sample complexity of agnostic learning under non-interactive LDP.

We also consider the problem of realizable learning and realizable refutation under non-interactive LDP, giving an algorithm which solves the realizable versions of learning and refutation and a lower bound which shows that this algorithm is nearly optimal for refutation.

Edmonds, A., A. Nikolov, J. Ullman (2020). "The Power of Factorization Mechanisms in Local and Central Differential Privacy". Symposium on Theory of Computing, STOC 2020. https://arxiv.org/abs/1911.08339.

This work gives a general characterization for answering linear queries under non-interactive local differential privacy (LDP). In particular, we provide a generalization of the factorization mechanism. Giving a lower bound, we show this approach to be optimal in sample complexity. By extension, we also obtain an algorithm which solves both agnostic learning and agnostic refutation under non-interactive LDP. By extension of our lower bound for linear query release, we show that this algorithm is optimal for refutation. Thus, refutation implies learnability under non-interactive LDP, while the converse remains open.

Edmonds, A. (2017). "Concepts of Efficient Samplability". MSc Thesis.

My MSc at the University of Toronto was supervised by Dan Roy, with whom I studied formalizations of the notion of computationally efficient probabilistic sampling, a concept regularly defined in widely varied ways. In particular, we focused on the distinction between approximate sampling to arbitrary precision versus perfect sampling. Our primary question was whether the latter could be obtained as a consequence of the former. Both the negative and positive answers to this question were studied in relation to traditional open problems in complexity theory, such as $P \stackrel{?}{=} NP$.

PROJECTS

CSC2401 Project: Unsupervised translation between dialects with Word2Vec 2016

This project explored techniques for unsupervised translation between dialects, namely languages whose vocabularies sufficiently overlap, by way of the distributed representation of words, Word2Vec, by Mikolov et al. These techniques were applied towards translation from both Middle English and Elizabethan English to contemporary English and vice versa.

CSC2515 Project: Gender classification with bi-directional recurrent neural networks 2015

This joint project with classmate Noah Fleming, applied an LSTM neural network towards classification of writings in teen fiction according to the gender of the author. Application of support vector machines to vectors produced by Word2Vec was also tested. With a training corpus of 40,000 documents, the rate of success on the test corpus was 83%.

Undergraduate Final Project: Gödel's incompleteness theorem 2014

As a final project for the Mathematics Specialist program at the University of Toronto, I gave a four-hour seminar for my peers which introduced important definitions and concepts from logic, and proceeded with the classic proof of Gödel's First Incompleteness Theorem.

PRESENTATIONS

Conference on Neural Information Processing Systems (NeurIPS) 2022
Learning versus Refutation in Noninteractive Local Differential Privacy

Symposium on Theory of Computing (STOC) 2020
The Power of Factorization Mechanisms in Local and Central Differential Privacy

Theory and Practice of Differential Privacy (poster presentation) 2019
The Power of Factorization Mechanisms in Local and Central Differential Privacy

Brookfield Institute for Innovation and Entrepreneurship 2019
Introduction to Differential Privacy

CONFERENCES & WORKSHOPS

Conference on Neural Information Processing Systems (NeurIPS) 2022

Symposium on the Theory of Computing (STOC) 2020

ACM CCS 2019 2019
Theory and Practice of Differential Privacy

Simons Institute for the Theory of Computing 2019
Privacy and the Science of Data Analysis;
Beyond Differential Privacy;

ACM CCS 2018 2018
Theory and Practice of Differential Privacy

Toronto Metropolitan University 2017
CanaDAM: Canadian Discrete and Algorithmic Mathematics Conference

Simons Institute for the Theory of Computing 2017
Computational Challenges in Machine Learning;
Representation Learning;
Interactive Learning;
Foundations of Machine Learning Boot Camp

Institute for Advanced Study 2016
Avi Wigderson is 60: A Celebration of Mathematics & Computer Science

Carnegie Mellon University 2015
PoCo: Summer School in Polyhedral Combinatorics

Queen Mary University 2015
Existential Polytime and Polyhedral Combinatorics:
Two minicourses by Jack Edmonds

REVIEW

Reviewer for:

- *ACM Transactions on Algorithms (TALG)* 2020
- *IEEE Symposium on Foundation of Computer Science (FOCS)* 2020
- *Journal of Privacy and Confidentiality* 2019

COURSES

A selection of courses taken at the University of Toronto.

Graduate

- CSC2556 *Algorithms for Collective Decision Making*
- CSC2429 *Proof Complexity, Mathematical Programming and Algorithms*
- CSC2501 *Computational Linguistics*
- CSC2506 *Probabilistic Graphical Models*
- CSC2515 *Machine Learning*
- CSC2404 *Computability and Logic*
- CSC2401 *Computational Complexity*

Undergraduate

- MAT357, MAT457, MAT458 *Real Analysis*
- MAT347 *Groups, Rings & Fields*
- MAT461 *Combinatorial Method*
- MAT332 *Complex Analysis*
- MAT327 *Topology*
- MAT309 *Mathematical Logic*

- MAT332 *Graph Theory*
- MAT402 *Classical Geometry*

TEACHING

Course Instructor – *University of Toronto* 2018

Responsible for all aspects of the course, including: delivering and planning lectures; organizing tutorials; creating assignments, tests, and marking schemes; providing office hours; supervising teaching assistants.

- CSC236 *Introduction to the Theory of Computation*

Graduate Teaching Assistant – *University of Toronto* 2015 - present

Led tutorials, held office hours, and graded assignments and exams.

- CSC236 *Introduction to the Theory of Computation*
- CSC240 *Enriched Introduction to the Theory of Computation*
- CSC263 *Data Structures and Analysis*
- CSC2412 *Algorithms for Private Data Analysis* (graduate course)
- CSC2412 *AI and Ethics: Mathematical Foundations and Algorithms* (graduate course)

Graduate Teaching Assistant – *Ryerson University* 2014 - 2015

Led tutorials and graded assignments and exams.

- MTH110 *Discrete Math I*
- MTH210 *Discrete Math II*
- MTH141 *Linear Algebra*
- MTH240 *Calculus II*

COMMUNITY

Organizer and system administrator of Compter Science Graduate Student Union Matrix/Element server for instant messaging – *U of T* 2020-2023

CS Theory Student Seminar organizer – *U of T* 2020

CS Theory Student Retreat organizer – *U of T* 2019

CS Graduate Student Board Games Club organizer – *U of T* 2017 - 2020

OTHER INTERESTS

Open-source software and hardware. Zettelkasten. Contact improvisational dance. Singing. Music. Classic movies.