

I am a PhD student applying for full-time positions in Responsible AI, with particular interests in fairness and privacy. I am excited to apply my theoretical research background and industry internship experience towards framing and improving issues of trustworthiness in the ML pipeline.

Education

Boston University 2019 - 2024 (expected)
Ph.D. Computer Science GPA: 3.9

Adviser: [Sofya Raskhodnikova](#), Research focus: privacy-preserving algorithms, fairness

Princeton University 2014 - 2018
B.A. Mathematics, Minor in Computer Science GPA: 3.8

Industry experience

Applied Scientist Intern, Amazon (New York) Sept 2023 - Dec 2023

- Hosted by Srinivasan Sengamedu in the People Experiences and Technology Central Science team.
- Designed an algorithm that provides label correction in the training dataset for fairer predictions in the presence of label bias. Showed superior performance of algorithm compared to prior work both theoretically and empirically on benchmark datasets (see [4]).

Applied Scientist Intern, Amazon Web Services (Sunnyvale) Sept 2022 - Dec 2022

- Hosted by Shiva Kasiviswanathan in the Causality team and in collaboration with Aaditya Ramdas (CMU).
- Designed the first algorithms with rigorous theoretical guarantees for differentially private conditional independence testing and evaluated their performance empirically (see [3]).

Analyst, Health Outcomes and Economics, Analysis Group (Boston) Aug 2018 - Aug 2019

- Conducted data aggregation and analysis (in R) for two studies on cost-benefit analysis of new cancer drugs, leading to two publications.
- Was the sole front-end developer (in R Shiny) for two client deliverable applications that automated cost-benefit analysis studies.

ML work

BatchNorm in tradeoff between private and robust deep learning, Course Project Spring 2022

- Investigated whether the addition of BatchNorm improves accuracy for deep learning models trained to be both private and robust to adversarial noise.
- Showed positive results for image classification tasks on MNIST and CIFAR-10, using PyTorch and the TensorFlow privacy package.

CNN for lung X-rays, Course Project Spring 2021

- Classified lung disease type based on X-ray images. Performed transfer learning using AlexNet (TensorFlow, Jupyter on GPU). Compared performance with and without weight freezing of the bottom layers. Used techniques such as regularization (L1-L2, dropout), data augmentation, to achieve final accuracy of 84% on balanced test set.

Skills

Programming Python, PyTorch, Pandas, Scikit-learn, TensorFlow, R
Teaching I am very passionate about mentoring young students and have volunteered with two education NGOs in the past (Community House, Tutoring Plus)

Publications (authors listed alphabetically)

- [1] “Counting Distinct Elements in the Turnstile Model with Differential Privacy under Continual Observation.” Palak Jain, Iden Kalemaj, Sofya Raskhodnikova, Satchit Sivakumar, Adam Smith. **NeurIPS**, 2023.
- [2] “Node-Differentially Private Estimation of the Number of Connected Components.” Iden Kalemaj, Sofya Raskhodnikova, Adam Smith, Charalampos Tsourakakis. **PODS**, 2023.
- [3] “Differentially Private Conditional Independence Testing.” Iden Kalemaj, Shiva Kasiviswanathan, Aaditya Ramdas. **AISTATS**, 2024.
- [4] “FairLift: Label Correction for Fairer Predictions in the Presence of Label Bias.” Iden Kalemaj, Hien Pham, Srinivasan Sengamedu. **In submission**, 2024.
- [5] “Isoperimetric Inequalities for Real-Valued Functions with Applications to Monotonicity Testing.” Hadley Black, Iden Kalemaj, Sofya Raskhodnikova. **ICALP**, 2023.
- [6] “Performative Prediction in a Stateful World.” Gavin Brown, Shlomi Hod, Iden Kalemaj. **AISTATS**, 2022.
- [7] “Sublinear-Time Computation in the Presence of Online Erasures.” Iden Kalemaj, Sofya Raskhodnikova, Nithin Varma. **ITCS**, 2022.