Sameer Wagh

Remote – Soda Hall, Berkeley, CA - 94720

+1 (xxx) xxx-7233 • swagh@alumni.princeton.edu • snwagh.github.io
sameer-wagh • snwagh

Research Interests

Security and Privacy: Broadly interested in (learning and) designing, building privacy-preserving systems. Current research interests include privacy-preserving machine learning, applied cryptography, and design of efficient cryptographic primitives as well as their implementation and deployment.

Education

University of California, Berkeley — **Berkeley, CA**
Post Doctoral Researcher — *2020–present*

Princeton University — **Princeton, NJ**
PhD in Electrical Engineering — *2014–2020*
Thesis: New Directions in Efficient Privacy-Preserving Machine Learning

Indian Institute of Technology, Madras — **TN, India**
BTech in Engineering Physics with Honors, CGPA – 9.11 (Core GPA – 9.39) — *2010–2014*
Minor in Mathematics

Honors and Awards

2020: Finalist for the Bede Liu Best Dissertation Award 2020.

2020: Recipient of the Facebook Systems for ML research award.

2019: Winner of the Qualcomm Innovation Fellowship, North America region.

2013: Best Solution Award for analytic mapping of stock prices to a 2D-Ising model for the GS Quantify Competition by Goldman Sachs.

2011: Among the top 4 students across India to receive Nurture program scholarship by NBHM.

2009: Recipient of the P.L. Bhatnagar Memorial Award, as the top scorer of the Indian team at IMO.

2009: Silver Medalist and India's highest scorer at the 50th International Maths Olympiad, Germany.

2008: Recipient of the prestigious Kishore Vaidyanik Protsahan Yojana (K.V.P.Y.) fellowship.

2008: Youngest student in India to qualify for the Regional Maths Olympiad and the Indian National Maths Olympiad.

Publications

- Eleftheria Makri*, Dragoș Rotaru*, Frederik Vercauteren*, Sameer Wagh*. "Rabbit: Efficient Comparison for Secure Multi-Party Computation."
  Financial Cryptography and Data Security (FC) 2021. Acceptance Rate: 25.3%
  [PDF] [Source Code] *author list alphabetical
- Sameer Wagh. "New Directions in Efficient Privacy-Preserving Machine Learning."
  PhD Dissertation 2020.
  [PDF]
- David Marco Sommer, Liwei Song, Sameer Wagh, and Prateek Mittal. "Towards Probabilistic Verification of Machine Unlearning."

Under submission.
[PDF] [Source Code]

- Sameer Wagh, Shruti Tople, Fabrice Benhamouda, Eyal Kushilevitz, Prateek Mittal, and Tal Rabin. "Falcon: Honest-Majority Maliciously Secure Framework for Private Deep Learning."
Privacy Enhancing Technologies Symposium (PETS) 2021. Acceptance Rate: 17%
[PDF] [Qualcomm Award] [Facebook Award] [News] [News] [Source Code]
- Hao Chen*, Miran Kim*, Ilya Razenshteyn*, Dragoș Rotaru*, Yongsoo Song*, and Sameer Wagh*. "Maliciously Secure Matrix Multiplication with Applications to Private Deep Learning."
International Conference on the Theory and Application of Cryptology and Information Security (AsiaCrypt) 2020. Acceptance Rate: 25.9%
[PDF] *author list alphabetical
- Sameer Wagh, Xi He, Ashwin Machanavajjhala, Prateek Mittal. "DP-Cryptography: Marrying Differential Privacy and Cryptography in Emerging Applications."
Communications of the ACM (CACM) 2020.
[PDF]
- Sameer Wagh, Divya Gupta, and Nishanth Chandran. "SecureNN: 3-Party Secure Computation for Neural Network Training."
Privacy Enhancing Technologies Symposium (PETS) 2019. Acceptance Rate: 22%
[PDF] [Source Code] [Deployment at Cape Privacy] [Deployment at OpenMined]
- Gerry Wan, Aaron Johnson, Ryan Wails, Sameer Wagh, and Prateek Mittal. "Guard Placement Attacks on Path Selection Algorithms for Tor."
Privacy Enhancing Technologies Symposium (PETS) 2019. Acceptance Rate: 22%
[PDF] [Source Code]
- Hans Hanley, Yixin Sun, Sameer Wagh, Prateek Mittal. "DPSelect: A Differential Privacy Based Guard Relay Selection Algorithm for Tor."
Privacy Enhancing Technologies Symposium (PETS) 2019. Acceptance Rate: 22%
[PDF]
- Sameer Wagh, Paul Cuff and Prateek Mittal. "Differentially Private Oblivious RAM."
Privacy Enhancing Technologies Symposium (PETS) 2018. Acceptance Rate: 16%
[PDF] [Source Code] [News]
- Manuel Costa, Lawrence Esswood, Olya Ohrimenko, Felix Schuster, and Sameer Wagh, "The Pyramid Scheme: Oblivious RAM for Trusted Processors."
Tech Report, 2017.
[PDF]
- Yanqi Zhou, Sameer Wagh, Prateek Mittal and David Wentzlaff, "Camouflage: Memory Traffic Shaping to Mitigate Timing Attacks."
International Symposium on High-Performance Computer Architecture (HPCA) 2017. Acceptance Rate: 22.3%
[PDF]

Patents

Private Deep Neural Network Training **February 2018**
Sameer Wagh, Divya Gupta and Nishanth Chandran
Patent number: 403629-US-PSP / SLW Ref: 1777.778PRV

Tunable Oblivious RAM **January 2015**
Sameer Wagh, Paul Cuff and Prateek Mittal
Patent number: US20170185534 A1 granted to Princeton University.

Work Experience

Research Internship at Microsoft Research, Redmond, USA **Summer 2019**
Privacy preserving analytics for machine learning

Internship in Applied MPC and Implementations, Bar Ilan University, Israel **Summer 2018**
Implementing efficient MPC primitives and protocols.

Research Internship at Microsoft Research, Bangalore, India **Summer 2017**
Developed efficient cryptographic protocols for privacy preserving machine learning.

Research Internship at Microsoft Research, Cambridge, UK **Summer 2016**
Efficient ORAM protocol implementation in a secure processor environment (SGX).

Research Assistant at Princeton University **Fall 2014**
Differentially Private Oblivious RAM protocol design and implementation

B.Tech Project at IIT Madras **2013-2014**
Quench dynamics across a first order transition in Ashkin Teller model.

Nurture Program by NBHM at TIFR, Bombay. **2011-2013**
BSc equivalent study in pure Mathematics (Algebra, Analysis and Topology)

Research Internship at Okinawa Institute of Science and Technology, Japan **Summer 2013**
Theoretical imaging of magnetic monopoles in frustrated spin-ice systems.

Research Internship at Indian Institute of Science Education and Research, Pune **Summer 2012**
Exploring magnetic traps to manipulate Bose Einstein Condensates.

Research Internship at Oneirix Labs, Pune, India **Summer 2012**
Efficient signal processing for noise cancellation.

Teaching and Mentoring

Teaching Assistant for ELE 535 **Fall 2015**
Teaching assistant for "Machine Learning and Pattern Recognition: Introduction to the theory and practice of machine learning."

Mentoring Senior Independent Work **2017-2019**
Mentoring senior independent work at Princeton for academic years 2017, 2018, and 2019.

- Gerry Wan: Winner of the Calvin Dodd MacCracken Senior Thesis award. Currently pursuing doctorate at Stanford University.
- Hans Hanley: Winner of the Daniel M. Sachs Class of 1960 Graduating Scholarship. Currently pursuing masters at Oxford University.

Prison Teaching Initiative **2017-2019**
Initiative to reduce incarceration rates by increasing access to post-secondary education. I've taught the following courses at the Garden State Youth Correctional Facility:

- Intermediate Algebra (MAT 030)
- Precalculus II (MAT 113)

PH101, MA101 Coaching Initiative **2013-2014**
Initiative to improve the performance of freshmen in introductory Math and Physics courses.

Dissertation Contributions

My dissertation focuses on 3 main contributions to the domain of privacy-preserving computation.

- SecureNN: A novel cross-layer protocol design paradigm for efficient non- linear operations in MPC. The approach avoids the use of expensive cryptographic techniques such as OTs and GCs in favor of simple modular arithmetic. This improves upon the prior art in private ML by about an order of magnitude. The simplicity of this approach has already led to its early adoption in the industry
- Falcon: This work improves upon SecureNN to further improve the performance of non-linear operations while operating in a stronger adversarial model. Protocols are secure against malicious adversaries (honest majority) and this is the first secure framework to support high capacity networks with over a hundred million parameters such as VGG16 as well as the first to support batch normalization, a critical component of deep learning that enables training of complex network architectures such as AlexNet.
- Ponytail: This work is the first demonstration of an $O(n^2)$ communication overhead matrix-multiplication of two $n \times n$ matrices in a dishonest majority MPC setting; the prior best known algorithm used an $O(n^3)$ communication. The work also makes a compelling case for the use of a hybrid approach – combining HE with MPC – for efficient privacy-preserving computation.

While improving the performance of privacy-preserving computation techniques, this dissertation also challenges research dogmas by demonstrating the elimination of inter-conversion protocols and demonstrating compute-bound MPC protocols. These foundational ideas have the potential to construct a new line of practical privacy-preserving protocols.

Research Contributions

I have also worked on a number of other research projects in the space of privacy technologies.

- Right to be forgotten, also known as the right to erasure, is the right of individuals to have their data erased from an entity storing it. While there exist a number of law and policy around this topic, there are very few concrete mechanisms whereby users can verify if service providers comply with their deletion requests. In the first of its kind, we propose formal framework to study the design of such verification mechanisms for data deletion requests – also known as machine unlearning – in the context of systems that provide machine learning as a service.
- Differential privacy (DP) has arisen as the state-of-the-art metric for quantifying individual privacy when sensitive data are analyzed. I have worked on a couple of projects at the intersection of differential privacy and cryptography. We explore if relaxed notions of cryptographic primitives can enhance their performance as well as look at formalizing the vast literature of differentially private primitives through a unique taxanomy.
- Tor is an open-source software for anonymous communication which has seen practical deployment. However, its importance and popularity has made it an attractive target for a number attacks – including from larger network level adversaries. While location-based path selection algorithms have been proposed as a countermeasure to defend against such attacks, their location-awareness can be used to exploit the system. We formalize the leakage as well as the space of such attacks in works such as DPSelect and guard placement attacks in Tor.

Developer Experience

Matrix-multiplication using fully homomorphic encryption (FHE) **2019**

Implementation of a fast, generic, and scalable private matrix-multiplication library over the production level Microsoft SEAL (here) codebase. This was a part of the first effort to generate matrix triples in a dishonest majority setting. The library introduces a number of optimizations for efficient implementation of fully homomorphic encryption. The library is implemented in `C++` and currently closed source.

Library for private deep learning using multi-party computation (MPC) **2017-2018**

Implementation of a private deep learning library with state-of-the-art protocols for private deep learning (here). The library supports dynamic addition of various neural network layers and can be deployed as-is over Amazon EC2 servers. This work has already seen industry adoption – DropoutLabs (here, here) over TensorFlow and at OpenMined (here) over PyTorch. The library is implemented in `C++` and is open sourced here.

Secure ORAM implementation over SGX **2017**

Implementation of the first fully secure Oblivious RAM implementation over Intel Software Guard Extensions (SGX). The library is secure against side channel attacks due to the use of hardware level oblivious operations. This work is a part of the confidential computing project (here) at Microsoft Research, Cambridge, UK. The library is implemented in C++ and is currently closed source.

Talks

गोपनीयता और संगणना (Privacy and Computation)

- Mahatma Gandhi Antarrashtriya Hindi VishwaVidyalaya — Feb 2021

Maliciously Secure Matrix Multiplication with Applications to Private Deep Learning

- Theory and Application of Cryptology and Information Security (AsiaCrypt) — Dec 2020

The Rise of Privacy Enhancing Technologies

- Microsoft Research. Redmond — Feb 2020
- AI Research Division, JP Morgan — Feb 2020
- Aarhus University — Nov 2019
- Katholieke Universiteit te Leuven (KU Leuven) — Nov 2019
- École Polytechnique Fédérale de Lausanne (EPFL) — Nov 2019
- RISE Lab, University of California, Berkeley — Oct 2019

Private Deep Learning Made Practical

- Qualcomm, San Diego — Oct 2019

SecureNN: 3-Party Secure Computation for Neural Network Training

- Facebook FAIR, New York — Feb 2019
- Google Deepmind, London — Oct 2018
- IBM TJ Watson Research Center — Sep 2018
- Privacy Enhancing Technologies Symposium, Barcelona — July 2018

Differentially Private Oblivious RAM

- Privacy Enhancing Technologies Symposium, Barcelona — July 2018

Understanding the Mysterious: Bitcoin

- INSPIRE Meetings, Electrical Engineering, Princeton — Jan 2016

Consensus and Byzantine Fault Tolerance

- GSS, Math Department, Princeton — Feb 2015

Path Integrals: Techniques and Applications; Quench Dynamics in the Ashkin Teller Model

- Boltzmann Club, IIT Madras — 2012-2014

Introduction to Groups, Group Representation, Character Theory and Applications in Physics

- Advanced Statistical Mechanics of Fields, IIT-Madras — Nov 2013

Fractals: A Measure Theoretic Introduction

- Advanced Dynamical Systems, IIT-Madras — Oct 2012

Other Services

PC Member

- Privacy Enhancing Technologies Symposium (PETS) 2022
- Privacy Enhancing Technologies Symposium (PETS) 2021
- Distributed & Privacy Preserving Machine Learning (ICLR Workshop) 2021
- Privacy-Preserving Machine Learning in Practice (CCS Workshop) 2020

Peer Reviewing

- USENIX Security Symposium (USENIX) 2016, 2017, 2018, 2019
- Privacy Enhancing Technologies Symposium (PETS) 2018, 2019, 2020
- IEEE Symposium on Security and Privacy (S&P) 2019, 2020, 2021
- Network and Distributed System Security Symposium (NDSS) 2017, 2018
- ACM Conference on Computer and Communications Security (CCS) 2021
- Communications of the ACM (CACM) 2020
- Practice and Theory in Public Key Cryptography (PKC) 2019, 2020
- Theory and Application of Cryptology and Information Security (AsiaCrypt) 2019

Organizing MelodEE 2017, 2018

Planning and organizing MelodEE, the annual talent show of ELE department, Princeton University.

Soccer Captain, Varsity Team 2012-2013

Led the university soccer team for the Inter IIT's, the annual sports tournament among all the IIT's. Other responsibilities include organizing Sports Fest (IIT Madras's annual sports tournament), Schroeter (Inter hostel tournaments) and all other university level tournaments.