Sikha Pentyala

Redmond, WA sikha@uw.edu +1 206 724 4643 sikhapentyala.github.io

Summary

Researcher with a focus on Responsible and Trustworthy AI, specializing in privacy-preserving techniques. My research spans privacy, synthetic data generation, fairness and interpretability. Dedicated to advancing AI in ethical and impactful ways through engaging with researchers from diverse backgrounds.

Education

Ph.D. in Computer Science and Systems Spring 2022 - Exp. Winter 2025
University of Washington Tacoma (GPA: 4.0/4.0)

Master of Science in Computer Science and Systems 2019 - 2020
University of Washington Tacoma (GPA: 4.0/4.0)

Post Graduate Diploma in Nuclear Engineering 2010 - 2011
Homi Bhabha National Institute (GPA: 4.0/4.0)

Bachelor of Computer Science 2005 - 2009
Jawaharlal Nehru Technological University, Anantapur (Gold Medalist - GPA: 4.0)

Experience

AI Research Summer Associate Jun 2024 – Aug 2024
JPMorgan Chase *New York, NY*
- Research on the interpretability of time series models

AI Research Summer Associate Jun 2023 – Aug 2023
JPMorgan Chase *New York, NY*
- Research on the intersection of privacy and explainability in ML

Pre-doctoral Instructor Dec 2022 – Jun 2023
University of Washington Tacoma *Tacoma, WA*
- Served twice as the primary instructor for Fundamentals of OOPs.

Research Intern Jul 2022 – Sep 2022
Mila - Quebec AI Institute *Remote*
- Research on privacy-preserving fair item rankings

Research Assistant (Privacy-preserving machine learning) Mar 2022 onwards
University of Washington Tacoma *Tacoma, WA*

Research Intern Oct 2021 – Mar 2022
Mila - Quebec AI Institute *Remote*
- Research on training machine learning models in a federated environment while preserving privacy of the users and achieving group fairness.

Graduate Research Student Jul 2020 - Oct 2020
University of Washington Tacoma - Microsoft *Tacoma, WA*
- Studied and designed system to automatically fix connectivity issues in existing Open Street Maps (OSM) dataset for its efficient use in routing services using C# in Visual Studio.
- System detected wrongly connected roadways to nodes for each country - 11% in New Zealand, 8% in the Fiji Islands, 19% in Venezuela and 8% in Serbia.

Scientific Officer Sep 2011 - Apr 2018
Nuclear Power Corporation of India Ltd. *Kudankulam, TN, India*
- Led a team of 2 developers to implement and re-design the intranet website to automate and achieve 50% faster process-lines, reducing labor and time costs.
- Estimated and managed cost, work scope, manpower, schedule and delivery for the intranet website while

leading a team of 2 software developers.

- Designed business specific applications for website such as online tender preparation, a complaint management system, a works and contracts management system and a chemistry data management system.
- Increased up-time by 20% by moving 10 legacy systems to new hardware and software while administering 14 server hardware as well as 27 software applications.
- Improved Annual Budget Exercises(Capital & Revenue) thru' web application leading to 70% faster delivery of budget and cost reports for each department and completion of discussion rounds for adhoc cuts.
- Performed market survey and procurement activities for IT infrastructure to purchase products at best offered price.

Teaching

Grader Fall 2024
University of Washington Tacoma *Tacoma, WA*
TCSS 343 (Design and Analysis of Algorithms)
TCSS 555 (Machine Learning)

Pre-doctoral Instructor Spring 2023
University of Washington Tacoma *Tacoma, WA*
TCSS 143 (Fundamentals of Object-Oriented Programming)

Pre-doctoral Instructor Winter 2023
University of Washington Tacoma *Tacoma, WA*
TCSS 143 (Fundamentals of Object-Oriented Programming)

Publications

Proceedings and Journals

- **S. Pentyala**, M. Pereira, M. De Cock. *CaPS: Collaborative and Private Synthetic Data Generation from Distributed Sources*, 41st International Conference on Machine Learning (ICML), 2024 (acceptance rate: 27.9%)
- R.J.M. Maia, D. Ray, **S. Pentyala**, R. Dowsley, M. De Cock, A. Nascimento and R. Jacobi. 2023. *An End-to-End Framework for Private DGA Detection as a Service*, PLOS ONE 2024 (acceptance rate: 20%)
- J. Vos, **S. Pentyala**, S. Golob, R. Maia, D. Kelly, C. Martins, Z. Erkin, M. De Cock, A. Nascimento. *Privacy-Preserving Membership Queries for Federated Anomaly Detection*, Proceedings on Privacy Enhancing Technologies (PoPETS), 2024 (acceptance rate: 20%)
- S. Golob, **S. Pentyala**, R. Dowsley, B. David, M. Larangeira, M. De Cock, A. Nascimento. *A Decentralized Information Marketplace Preserving Input and Output Privacy*, Proceedings of the 2nd Data Economy Workshop (DEC23) - SIGMOD 2023 Workshop, 2023
- Jia Ao Sun, **S. Pentyala**, M. De Cock and G. Farnadi. *Privacy-Preserving Fair Item Ranking*, 45th European Conference on Information Retrieval 2023 (acceptance rate: 29%)
- **S. Pentyala**, N. Neophytou, A. Nascimento, M. De Cock, G. Farnadi *Privacy-Preserving Group Fairness in Cross-Device Federated Learning*, Proceedings of Algorithmic Fairness through the Lens of Causality and Privacy (AFCP2022)
- **S. Pentyala**, R. Dowsley and M. De Cock *Privacy-Preserving Video Classification with Convolutional Neural Networks*, 38th International Conference on Machine Learning (ICML), 2021 (acceptance rate: 21.46%)
- F. Tabet and **S. Pentyala**, B. Patel, et al. *OSMRunner: A System for Exploring and Fixing OSM Connectivity*, 22nd IEEE International Conference on Mobile Data Management (MDM), 2021

Workshops and Abstracts

- T. Claar, S. Golob, **S. Pentyala**, G. Sitaraman, M. De Cock, J. Banerjee, L. Foschini. *Securely Generating Synthetic Genomic Data from Distributed Data Silos*, 11th International Workshop on Genome Privacy and Security (GenoPri'24), 2024
- S. Golob, **S. Pentyala**, A. Maratkhan, M. De Cock. *High Epsilon Synthetic Data Vulnerabilities in MST and PrivBayes*, PPAI-2024 (AAAI-24 Workshop on Privacy-Preserving Artificial Intelligence), 8 pages

- **S. Pentyala**, S. Sharma, S. Kariyappa, F. Leuce and D. Magazzeni. *Privacy-Preserving Algorithmic Recourse*, 3rd International Workshop on Explainable AI in Finance, ICAIF 2023, 8 pages
- M. Pereira, **S. Pentyala**, A. Nascimento, R. T. de Sousa Jr., M. De Cock *Secure Multiparty Computation for Synthetic Data Generation from Distributed Data*, SyntheticData4ML Workshop@NeurIPS 2022, 6 pages
- **S. Pentyala**, D. Melanson, M. De Cock and G. Farnadi. *PrivFair: a Library for Privacy-Preserving Fairness Auditing*, PPAI-2022 (AAAI-22 Workshop on Privacy-Preserving Artificial Intelligence), 8 pages
- **S. Pentyala**, N. Neophytou, A. Nascimento, M. De Cock, G. Farnadi . *PrivFairFL: Privacy-Preserving Group Fairness in Federated Learning*, Montreal AI Symposium (MAIS) 2022
- **S. Pentyala**, N. Neophytou, A. Nascimento, M. De Cock, G. Farnadi . *Towards private and fair federated learning*, WiML@NeurIPS 2022 (abstract/poster in Women in Machine Learning Workshop at NeurIPS, 2022)
- **S. Pentyala**, M. De Cock and R. Dowsley *Privacy-Preserving Video Classification*, WiCV 2021 (extended abstract/poster in Women in Computer Vision Workshop at CVPR, 2021)

Patents

- System and method for generating recourse paths with privacy guarantees (Application number 18/517,211)
- Pending undisclosed patent on Time Series at JPMorgan Chase

Services

- Invited talk on privacy-preserving AI across data silos at RAISE 2024, UW Seattle
- Contributing talk on synthetic data from distributed silos at WiDS UW Tacoma 2024
- Contributed blog posts on Privacy-Preserving Federated Learning Blog Series
- Co-organizing the WiDS table at Sisterhood Event at UW SET Tacoma
- Reviewer for ECAI 2024
- Volunteer at ICML 2024
- Reviewer for SyntheticData4ML Workshop@NeurIPS 2023, SyntheticData4ML Workshop@NeurIPS 2022.
- Reviewer for IEEE Transactions on Dependable and Secure Computing 2022
- Reviewer for IEEE Transactions on Neural Networks and Learning Systems 2022, 2023, 2024.
- Reviewer for International Journal of Information Management 2024
- Invited talk on wining solution to U.S. PETs Prize challenge at WiDS UW Tacoma 2023
- Contributed discussions to the invite-only virtual workshop on Advancing Privacy and Fairness in May 2023
- Volunteer in Women in Machine Learning @NeurIPS 2022
- Member of Women in Data Science
- Roundtable Lead on Privacy at the Algorithmic Fairness through the Lens of Causality and Privacy (AFCP2022) Workshop, NeurIPS 2022.
- Co-Reviewer for NeurIPS 2022

Awards

- 2023-2024 School of Engineering & Technology's Outstanding Student Leadership Award
- 2023 UPE Executive Council Award by UPE International Honor Society for the Computing and Information Disciplines
- 2023 JP Morgan Chase Fellowship to support PhD research on synthetic data generation
- SNAKE Challenge: won 1^{st} place on membership attacks against synthetically generation data.
- 2023 Andrew and Julie Fry Innovation Award, UW Tacoma
- 2022-2023 School of Engineering & Technology's Student Justice Equity Diversity & Inclusion Award, UW Tacoma
- 2022-2023 School of Engineering & Technology's Outstanding Graduate Research Award, UW Tacoma
- U.S.-U.K. PETs Prize Challenges: 2^{nd} prize in Phase 2 of Track A
- WiML travel funding award (WiML@NeurIPS 2022), PPAI travel funding award (PPAI@AAAI 2024), ICML travel funding award (ICML 2024)

- Carwein-Andrews Endowment award for 2022-2023, UW Tacoma, School of Engineering and Technology
- 2021-2022 School of Engineering & Technology's Outstanding Graduate Research Award, UW Tacoma
- Winner of Track III of the iDASH2021 secure genome analysis competition
- Gold Medalist at Jawaharlal Nehru Technological University, Anantapur in the Branch of ECM, Class of 2009

Relevant Projects

- **Privacy-preserving Generative AI - Privacy and Synthetic Data**
  - Built first solution to generate synthetic data from distributed silos for arbitrary partition - based on Secure Multiparty Computation (MPC) and Differential Privacy (DP).
  - Ongoing research on adapting the above solution for real-world application - secure genomic synthetic data generation.
- **Interpretable Models for Time Series - Explainability**
  - Research on interpretability of time series models.
  - Devised a new architecture for interpretable of time series models at feature and temporal level
- **Privacy-preserving Algorithmic Recourse - Privacy and Explainability**
  - Devised techniques based on Differential Privacy (DP) to provide multi-step recourse to the customers (end-users) while preserving the privacy of the individuals in the training dataset.
- **Attacks on Privacy-preserving Synthetic Data Generation - Privacy and Synthetic Data**
  - Member of the team stevengolob in the SNAKE Challenge.
  - Devised custom and novel techniques to perform membership attacks against marginal based synthetic data generators.
- **Private Financial Crime Detection - Privacy**
  - Member of the team PPMLHuskies in the U.S. PETs Prize Challenge.
  - Project on designing privacy-preserving federated learning solution to train models on vertically and horizontally partitioned data to detect anomalies in financial transactions.
- **Recommendation system - Privacy and Fairness**
  - Devised techniques based on secure multi-party computation (MPC) to preserve consumer (user) privacy while achieving fairness for producers (items).
- **Private and Fair ML training - Federated learning, Privacy and Fairness**
  - Devised techniques to achieve group fairness in federated learning (FL) while preserving privacy of the users and their sensitive attributes.
  - Developed MPC protocols for existing central pre-processing and post-processing techniques to adapt in a FL setup in privacy-preserving way.
  - Evaluated proposed techniques using Flower framework with 100+ clients.
- **Privacy preserving fairness auditing - Privacy and Fairness**
  - Built first MPC protocols to compute statistical notions of group fairness such as equalized odds, demographic parity and subgroup accuracy for multi-class and binary classification in MP-SPDZ.
  - Developed and deployed a demo web-based service that takes images or tabular data from one user and model parameters from the other user and performs secure auditing of the model using above MPC protocols. Audited a CNN with 1.48 parameters with 56 images in 30 secs in a 3PC passive setting.
- **Confidential Computing - Federated learning, MPC and Differential privacy**
  - Built a solution where two parties individually train a DP logistic regression (LR) model and obtain aggregated ϵ-DP LR model. Achieved a test accuracy of 85.79% with the model trained in 0.27s in the iDASH 2021 competition.
  - Built a solution where multiple parties collaborate to train a DP LR model using MPC, suitable for multiple scenarios of data splits. Achieved a test accuracy of 87.98% in 76s in a 3PC passive setting.
- **Privately classifying a video using Convolutional Neural Networks - Privacy with MPC**
  - Built first end-to-end MPC protocol for private video classification in Python using MP-SPDZ benchmarking software and Tensorflow.
  - Designed and trained a CNN with 1.5 million paramters to detect 7 emotions in a video from RAVDESS

dataset.
  - Deployed solution in cloud (AWS, Azure) & classified 4s video in 8.5s (passive 3PC) while keeping content of video and model parameters private to the owners in different adversarial models.

Skills

- Programming Languages: Java, Python, C#
- Frameworks: MP-SPDZ, FairLearn, Flower, Darts
- Tools and Technologies: Keras, Sklearn, Pandas, TensorFlow, Pytorch, OpenCV, Docker, Git (Github), Jupyter
- Databases: MS-SQL, ORACLE DB, MySQL
- Cloud Technologies: Amazon Web Services (AWS - EC2, S3, SageMaker), Azure Containers
- Servers: Apache-Tomcat, WAMP/LAMP, MS AD-DC and Exchange, BlueCoat