MIT-IBM Watson AI Lab
IBM Research
Cambridge, MA 02142, USA

Mobile: (+1)-315-744-6778
Email: lsjxjtu@gmail.com, sijia.liu@ibm.com
Website: http://sliu17.mysite.syr.edu/

PRIMARY RESEARCH AREAS

Machine learning: deep learning, adversarial ML, explainability, sparse learning, automated ML, computer vision
Optimization: black-box optimization, non-convex optimization, optimization for deep learning, distributed learning
Data science: time-series data analysis, network data analysis, chemical and biological data analysis
Signal processing: estimation/detection, graph signal process, information fusion
Computational biology: genome architecture and transcription, system biology, cell reprogramming

WORK EXPERIENCE

Research Staff Member, MIT-IBM Watson AI Lab, IBM Research Jan. 2018 – Present

Postdoc Research Fellow, University of Michigan, Ann Arbor, MI April 2016 – Dec. 2017
Supervisors: Alfred Hero (EECS) and Indika Rajapakse (Computational Medicine & Bioinformatics)

Data Science PhD Intern, Huawei R&D USA, Bridgewater, NJ June 2015 – Aug. 2015

Research Assistant, Syracuse University, Syracuse, NY Sept. 2011 – Mar. 2016
Advisors: Pramod K. Varshney (EECS) and Makan Fardad (EECS)

EDUCATION

Ph.D. in Electrical and Computer Engineering, Syracuse University Mar. 2016
Thesis: "Resource management for distributed estimation via sparsity-promoting regularization"
(**All University Doctoral Prize**)

M.S. in Electrical Engineering, Xi'an Jiaotong University May 2011

B.S. in Electrical Engineering, Xi'an Jiaotong University May 2008

AWARDS AND RECOGNITIONS

- **IBM Outstanding Research Accomplishments**, 2019
  — *Trustworthy AI; Toward Automating the AI Lifecycle with AutoAI; Deep Learning on Graphs*
- **Winner of Best Student Paper Award (3rd place)**, the 42nd IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2017
- **Recipient of All University Doctoral Prize**, Syracuse University, 2016
- **Best Student Paper Nominee** (among the seven finalists) at Asilomar Conference on Signals, Systems, and Computers, CA, Pacific Grove, CA, 2013
- **Winner of Best Poster Award** at Nunan Poster Competition, Syracuse University, 2012
- **First Class Award in National Mathematics Olympiad**, 2004
  — *Exempted from National College Entrance Examination in China*

GRANT EXPERIENCE

- **IBM PI**, *"Toward Trustworthy AI: Efficient Algorithms for Building Provably Robust and Verifiable Neural Networks"*, MIT-IBM AI Challenge Award, $750K, 2018 – 2021 (MIT PI Luca Daniel)
- **IBM PI**, *"Instruction, Command Line or Script Malware Detection"*, MIT-IBM AI Challenge Award, $750K, 2019 – 2022 (MIT PI Una-May O'Reilly)
- **IBM PI**, *"Fast Learning of Neural Network Models with Provable Generalizability"*, RPI-IBM AI Challenge Award, $150K, 2019 – 2020 (RPI PI Meng Wang)

SELECTED PUBLICATIONS

***Full* publications can be found at Google Scholar**

$*$ denotes equal contribution, † denotes first authors under my supervision/co-supervision.

AI/Machine learning

[1] A. Boopathy†, **S. Liu**, G. Zhang, P.-Y. Chen, S. Chang, and L. Daniel, "Visual Interpretability Alone Helps Adversarial Robustness", https://openreview.net/pdf?id=Hyes70EYDB

[2] **S. Liu***, S. Lu*, X. Chen*†, Y. Feng*†, K. Xu*†, A. Al-Dujaili*, M. Hong, and U.-M. Obelilly, "Min-Max Optimization without Gradients: Convergence and Applications to Adversarial ML", https://arxiv.org/abs/1909.13806

[3] M. Cheng, S. Singh, P.-Y. Chen, **S. Liu**, and C.-J. Hsieh, "Sign-OPT: A Query-Efficient Hard-label Adversarial Attack ", *International Conference on Learning Representations (ICLR'20)* (acceptance rate 26.5%)

[4] **S. Liu***, P. Ram*, D. Vijaykeerthy, D. Bouneffouf, G. Bramble, H. Samulowitz, D. Wang, A. R. Conn, and A. Gray "An ADMM Based Framework for AutoML Pipeline Configuration", *34th AAAI Conference on Artificial Intelligence (AAAI'20)* (acceptance rate 20.6%)

[5] P. Zhao*†, L. Weng*, **S. Liu**, P.-Y. Chen, X. Lin, and L. Daniel, "Towards Certificated Model Robustness Against Weight Perturbations", *AAAI'20* (acceptance rate 20.6%)

[6] **S. Liu***, X. Chen*†, K. Xu*†, X. Li*, X. Lin, M. Hong, and D. Cox, "ZO-AdaMM: Zeroth-Order Adaptive Momentum Method for Black-Box Optimization", *33rd Conference on Neural Information Processing System (NeurIPS'19)* (acceptance rate 21.1%)

[7] K. Xu*†, H. Chen*†, **S. Liu**, P.-Y. Chen, T.-W. Wen, M. Hong, and X. Lin, "Topology Attack and Defense for Graph Neural Networks: An Optimization Perspective", *International Joint Conference on Artificial Intelligence (IJCAI'19)* (acceptance rate 17.9%)

[8] P. Zhao†, **S. Liu**, P.-Y. Chen, N. Hoang, K. Xu, S. Wang, Y. Wang, and X. Lin, "On the Design of Black-box Adversarial Examples by Leveraging Gradient-free Optimization and Operator Splitting Method", *IEEE International Conference on Computer Vision 2019 (ICCV'19)* (acceptance rate 25%)

[9] S. Ye*†, K. Xu*†, **S. Liu**, H. Cheng, J.-H. Lambrechts, H. Zhang, A. Zhou, K. Ma, Y. Wang, and X. Lin, "Adversarial Robustness vs. Model Compression, or Both?", *ICCV'19* (acceptance rate 25%)

[10] T. Zhang†, **S. Liu**, Y. Wang, and M. Fardad, "Generation of Low Distortion Adversarial Attacks via Convex Programming", *IEEE International Conference on Data Mining (ICDM'19)* (acceptance rate 18.5%)

[11] P.-Y. Chen, L. Wu, **S. Liu**, I. Rajapakse, "Fast Incremental von Neumann Graph Entropy Computation: Theory, Algorithm, and Applications", *International Conference on Machine Learning (ICML'19)* (acceptance rate 22.6%)

[12] **S. Liu**, P.-Y. Chen, X. Chen, M. Hong, "signSGD via Zeroth-Order Oracle", *ICLR'19* (acceptance rate 31.4%)

[13] **S. Liu***, K. Xu*†, P. Zhao, P.-Y. Chen, H. Zhang, Q. Fan, D. Erdogmus, Y. Wang, and X. Lin "Structured Adversarial Attack: Towards General Implementation and Better Interpretability", *ICLR'19* (acceptance rate 31.4%)

[14] X. Chen†, **S. Liu**, R. Sun, and M. Hong. "On the Convergence of A Class of Adam-Type Algorithms for Non-Convex Optimization", *ICLR'19* (acceptance rate 31.4%)

[15] A. Boopathy†, L. Weng, P.-Y. Chen, **S. Liu**, and L. Daniel, "CNN-Cert: An Efficient Framework for Certifying Robustness of Convolutional Neural Networks", *AAAI'19* (acceptance rate 16.2%) (oral)

[16] C.-C. Tu*, P.-S. Ting*, P.-Y. Chen*, **S. Liu**, H. Zhang, J. Yi, C.-J. Hsieh, and S.-M. Chen, "AutoZOOM: Autoencoder-based Zeroth Order Optimization Method for Attacking Black-box Neural Networks", *AAAI'19* (acceptance rate 16.2%) (oral)

[17] **S. Liu**, B. Kailkhura, P.-Y. Chen, P. Ting, S. Chang and L. Amini, "Zeroth-Order Stochastic Variance Reduction for Nonconvex Optimization", *NeurIPS'18* (acceptance rate 22.7%).

[18] P. Zhao, **S. Liu**, Y. Wang, X. Lin, "An ADMM-Based Universal Framework for Adversarial Attacks on Deep Neural Networks", *ACM Multimedia (ACMMM)*, 2018

[19] **S. Liu**, J. Chen, P.-Y. Chen and A. O. Hero, "Zeroth-Order Online Alternating Direction Method of Multipliers: Convergence Analysis and Applications", *AISTATS'18*

[20] **S. Liu**, Y. Wang, M. Fardad and P. K. Varshney, "A Memristor-Based Optimization Framework for Artificial Intelligence Applications", *IEEE Circuits and Systems Magazine*, 2018

Computational biology

[21] **S. Liu**, H. Chen, S. Ronquist, L. Seaman, N. Ceglia, W. Meixner, L. A. Muir, P.-Y. Chen, G. Higgins, P. Baldi, S. Smale, A. O. Hero and I. Rajapakse, "Genome Architecture Mediates Transcriptional Control of Human Myogenic Reprogramming," *iScience, Cell*, 2018

[22] H. Chen, L. Seaman, **S. Liu**, T. Ried, and I. Rajapakse, "Chromosome conformation and gene expression patterns differ profoundly in human fibroblasts grown in spheroids versus monolayers," *Nucleus*, 2017

[23] H. T. Ali†, **S. Liu**, Y. Yilmaz, R. Couillet, I. Rajapakse, A. Hero, "Latent Heterogeneous Multilayer Community Detection", *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2019

Signal processing

[24] S. Zhang†, **S. Liu**, V. Sharma and P. K. Varshney, "Optimal Sensor Collaboration for Parameter Tracking Using Energy Harvesting Sensors", *IEEE Trans. Signal Process.*, 2018

[25] **S. Liu**, P.-Y. Chen and A. O. Hero, "Accelerated Distributed Optimization for Evolving Networks of Growing Connectivity", *IEEE Trans. Signal Process.*, 2017

[26] **S. Liu**, S. Kar, M. Fardad and P. K. Varshney, "Optimized Sensor Collaboration for Estimation of Temporally Correlated Parameters", *IEEE Trans. Signal Process.*, 2016

[27] **S. Liu**, S. P. Chepuri, M. Fardad, E. Masazade, G. Leus and P. K. Varshney, "Sensor Selection for Estimation with Correlated Measurement Noise", *IEEE Trans. Signal Process.*, 2016

[28] B. Kailkhura, **S. Liu**, T. Wimalajeewa and P. K. Varshney, "Measurement Matrix Design for Compressive Detection with Secrecy Guarantees", *IEEE Wireless Commun. Lett.*, 2016

[29] **S. Liu**, S. Kar, M. Fardad and P. K. Varshney, "Sparsity-Aware Sensor Collaboration for Linear Coherent Estimation", *IEEE Trans. Signal Process.*, 2015

[30] **S. Liu**, A. Vempaty, M. Fardad, E. Masazade and P. K. Varshney, "Energy-Aware Sensor Selection in Field Reconstruction", *IEEE Signal Process. Lett.*, 2014

[31] X. Shen, **S. Liu** and P. K. Varshney, "Sensor Selection for Nonlinear Systems in Large Sensor Networks", *IEEE Trans. Aerosp. Electron. Syst.*, 2014

[32] **S. Liu**, M. Fardad, E. Masazade and P. K. Varshney, "Optimal Periodic Sensor Scheduling in Large-Scale Dynamical Networks", *IEEE Trans. Signal Process.*, 2014

[33] P.-Y. Chen and **S. Liu**, "Bias-Variance Tradeoff of Graph Laplacian Smoothing Regularizer", *IEEE Signal Process. Lett.*, 2017

[34] **S. Liu**, A. Ren†, Y. Wang and P. K. Varshney, "Ultra-Fast Robust Compressive Sensing Based on Memristor Crossbars," *ICASSP*, 2017 (Winner of Best Student Paper Award, 3rd place)

PRESS COVERAGE

- **VentureBeat :** *Researchers foil people-detecting AI with an 'adversarial' T-shirt* October 2019
- **IBM Research Blog:** *Making Neural Networks Robust with New Perspectives* August 2019
- **Medium:** *AI Safety - How Do you Prevent Adversarial Attacks?* August 2019
- **IBM Research Blog:** *Will Adam Algorithms Work for Me?* May 2019
- **Medium:** *CNN-Cert: A Certified Measure of Robustness for Convolutional Neural Networks* January 2019
- **IBM Research Blog:** *Efficient Adversarial Robustness Evaluation of AI Models with Limited Access* January 2019

SERVICE

- **Co-chair** of IBM AI Research Week Workshop *Foundations of Safe Learning*, 2019
- **Co-chair** of KDD Workshop *Adversarial Learning Methods for Machine Learning and Data Mining*, 2019
- **Co-chair** of IEEE GlobalSIP Workshop *Signal Processing for Adversarial Machine Learning*, 2018
- **Co-chair** of ICME workshop *Machine Learning and Artificial Intelligence for Multimedia Creation*, 2018
- **Guest editor**, *IEEE Internet of Things Journal special issue on AI Enabled Cognitive Communications and Networking for IoT*, 2018
- **Vice-chair** of *IEEE ComSoc SIG on AI Embedded Cognitive Networks*, 2017-present
- **Referee for journals**: *IEEE Transactions on Information Theory, IEEE Transactions on Signal Processing, IEEE Transactions on Wireless Communications, IEEE Transactions on Automatic Control, Information Fusion, IFAC Journal of Automatica, IEEE Internet of Things Journal, IEEE Sensors Journal, etc*
- **Program committee member for conferences**: *NeurIPS, ICML, ICLR, AAAI, CVPR, ICCV, UAI, IJCAI, ACMMMM, ICASSP, GlobalSIP, CDC, ACC*