

Azalia Mirhoseini

CONTACT INFORMATION 1600 Amphitheatre Parkway
Mountain View, CA 94043 azaliamirh@gmail.com
azalia@google.com

TWITTER <https://twitter.com/azaliamirh>

WEBSITE www.azaliamirhoseini.com

RESEARCH AREA Deep Reinforcement Learning, Metalearning, Scalable Deep Learning

RECENT HIGHLIGHTS

- Co-founder/Leader of Machine Learning for Systems Moonshot at Google Brain, a larger research effort to use machine learning to design and optimize computer systems
- Developed deep reinforcement learning algorithms to do model parallelism that speed up deep network training by more than 60% over top-performing baselines (ICLR'17 and ICLR'18)
- Developed NLP architectures that allow efficient training of massive neural networks with over 100 billion parameters, yielding state of the art results on established language modeling and translations tasks (ICLR'17, > 200 citations)
- Developed new state-of-the-art Vision models on CIFAR-10 (2.01%) and CIFAR-100 (12.40%) by designing new search spaces and inserting ensembling into Neural Architecture Search Framework
- Mentor of OpenAI 2019 Scholars Program
- Co-organized "ML for Systems" workshop with > 300 attendees at NeurIPS 2018
- Co-chair of ML Day at International SuperComputing Conference 2019
- Keynote speaker, Invited Speaker, and Panelist at workshops/events with several hundreds of attendees at NeurIPS, SuperComputing Conference, Tsinghua AI Summit, and Frontiers of AI

PROFESSIONAL EXPERIENCE

Senior Research Scientist
Google Brain, Mountain View, CA 2016-Present

Postdoctoral Researcher
Department of Electrical and Computer Engineering
University of California, San Diego and Rice University, Houston, TX 2015-2016

Graduate Researcher
Department of Electrical and Computer Engineering
Rice University, Houston, TX 2009-2015

Research Intern
Sensing and Energy Research Group
Microsoft Research, Redmond, WA 2013

EDUCATION

Rice University, Houston, TX

Ph.D., Electrical and Computer Engineering, May 2015

- Thesis: *A Data- and Platform-Aware Framework for Large-Scale Machine Learning*
- Advisor: *Prof. Farinaz Koushanfar*

Sharif University of Technology, Tehran, Iran

B.S., Electrical Engineering, May 2009

HONORS AND
AWARDS

- William Marsh Rice’s Best Thesis Award, ECE Department, 2015
- Schlumberger Fellowship, 2012-2013
- IBM Ph.D. Scholarship, 2012
- Microsoft Research Graduate Women’s Scholarship, 2010
- National Math Olympiad Gold Medal, Iran, 2004

CONFERENCE AND
JOURNAL
PUBLICATIONS

1. **A. Mirhoseini***, A. Goldie*, H. Pham, B. Steiner, Q. V. Le, and J. Dean. “Hierarchical Planning for Device Placement.” *International Conference on Learning Representations (ICLR)*, 2018
2. **A. Mirhoseini***, H. Pham*, Q. V. Le, B. Steiner, R. Larsen, Y. Zhou, N. Kumar, M. Norouzi, S. Bengio, and J. Dean. “Device Placement Optimization with Reinforcement Learning.” *International Conference on Machine Learning (ICML)*, 2017
3. [C]¹ N. Shazeer, **A. Mirhoseini***², K. Maziarz*, A. Davis, Q. V. Le, G. Hinton, and J. Dean. “Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer.” *International Conference on Learning Representations (ICLR)*, 2017
4. [J] **A. Mirhoseini**, E. Dyer, E. Songhori, R. Baraniuk, and F. Koushanfar. “RankMap: A Framework for Distributed Learning from Dense Datasets”. *IEEE Transactions on Neural Networks and Learning Systems (TNNLS)*, 2017
5. [C] B. Rouhani, **A. Mirhoseini**, and F. Koushanfar. “Deep³: Leveraging Three Levels of Parallelism for Efficient Deep Learning.” *Design Automation Conference (DAC)*, 2017
6. [J] B. Rouhani, **A. Mirhoseini**, and F. Koushanfar. “RISE: An Automated Framework for Real-Time Intelligent Video Surveillance on FPGA”. *ACM Transactions on Embedded Computing Systems (TECS)*, 2017
7. [C] B. Rouhani, **A. Mirhoseini**, and F. Koushanfar. “TinyDL: Just-in-time deep learning solution for constrained embedded systems.” *International Symposium On Circuits and Systems (ISCAS)*, 2017
8. [C] **A. Mirhoseini**, B. Rouhani, E. Songhori, and F. Koushanfar. “Perform-ML: Performance Optimized Machine Learning by Platform and Content Aware Customization.” *Design Automation Conference (DAC)*, 2016
9. [C] B. Rouhani, **A. Mirhoseini**, and F. Koushanfar. “DeLight: Adding Energy Dimension To Deep Neural Networks”. *International Symposium on Low Power Electronics and Design (ISLPED)*, 2016
10. [J] B. Rouhani, **A. Mirhoseini**, E. Songhori, and F. Koushanfar. “Automated Analysis of Streaming Big and Dense Data on Reconfigurable Platforms.” *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, 2016
11. [C] B. Rouhani, **A. Mirhoseini**, and F. Koushanfar. “Going Deeper than Deep Learning for Massive Data Analytics under Physical Constraints.” *International Conference on Hardware/Software Co-design and System Synthesis (CODES+ISSS)*, 2016
12. [J] **A. Mirhoseini**, B. Rouhani, E. Songhori, and F. Koushanfar. “Chime: Checkpointing Long Computations on Intermittently Energized IoT Devices.” *IEEE Transactions on Multi-Scale Computing Systems (TMSCS)*, 2016

¹Conference papers are marked by [C] and journal papers are marked by [J].

²Major Contributor.

13. [C] **A. Mirhoseini**, A. Sadeghi, and F. Koushanfar. “Scalable and privacy-preserving outsourcing of learning algorithms.” *International Symposium on Hardware Oriented Security and Trust (HOST)*, 2016
14. [C] R. Patel, T. Goldstein, E. Dyer, **A. Mirhoseini**, and R. Baraniuk. “Adaptive column sampling and Nystrom approximation via oASIS.” *Siam International Conference on Data Mining (SDM)*, 2016
15. [C] **A. Mirhoseini** and F. Koushanfar. “Enabling privacy preserving computing at scale by modular signal processing.”, *Allerton Conference on Communication, Control, and Computing (Allerton)*, 2015
16. [C] **A. Mirhoseini**, E. Songhori, B. Rouhani, and F, Koushanfar. “Extensible Dictionaries for Customized Learning of Big Data.” *Short Paper. ACM Special Interest Group for the Computer Systems Performance Evaluation Conference (SIGMETRICS)*, 2015
17. [C] F. Koushanfar, **A. Mirhoseini**, G. Qu, and Z. Zhang. “DA Systemization of Knowledge: A Catalog of Prior Forward-Looking Initiatives.”, *International Conference on Computer-Aided Design (ICCAD)*, 2015
18. [C] B. Rouhani, E. Songhori, **A. Mirhoseini**, and F, Koushanfar. “SSketch: An Automated Framework for Streaming Sketch-based Analysis of Big Data on FPGA.” *IEEE International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, 2015
19. [C] E. Songhori, **A. Mirhoseini**, X. Lu, and F. Koushanfar. “AHEAD: Automated Framework for Hardware Accelerated Iterative Data Analysis.” *Design Automation and Test in Europe Conference (DATE)*, 2015
20. [J] **A. Mirhoseini**, M. Potkonjak, and F. Koushanfar. “Phase Change Memory Write Cost Minimization by Data Encoding.” *IEEE Journal on Emerging and Selected Topics in Circuits and Systems (JETCAS)*, 2015
21. [C] **A. Mirhoseini**, E. Songhori, and F. Koushanfar. “Automated Checkpointing for Enabling Intensive Applications on Energy Harvesting Devices.” *International Symposium on Low Power Electronics and Design (ISLPED)*, 2013
22. [C] **A. Mirhoseini**, E. Songhori, and F. Koushanfar. “Idetic: A High-level Synthesis Approach for Enabling Long Computations on Transiently-powered ASICs.” *IEEE International Conference on Pervasive Computing and Communications (PerCom)*, 2013
23. [C] **A. Mirhoseini**, M. Potkonjak, and F. Koushanfar. “Coding-Based Energy Minimization for Phase Change Memory.” *Design Automation Conference (DAC)*, 2012
24. [C] **A. Mirhoseini** and F. Koushanfar. “Learning to Manage Combined Energy Supply Systems.” *International Symposium on Low Power Electronics and Design (ISLPED)*, 2011
25. [C] **A. Mirhoseini** and F. Koushanfar. “HypoEnergy: Hybrid Supercapacitor-Battery Power-supply Optimization for Energy Efficiency.” *Design Automation and Test in Europe conference (DATE)*, 2011
26. [C] F. Koushanfar and **A. Mirhoseini**. “Hybrid Heterogeneous Energy Supply Networks.”, *International Symposium on Circuits and Systems (ISCAS)*, 2011

27. [J] F. Koushanfar and **A. Mirhoseini (Alphabetical Order)**. “A Unified Framework for Submodular Integrated Circuit Trojan Detection.” *IEEE Transactions on Forensics and Security (TIFS)*, 2011
28. [C] **A. Mirhoseini**, Y. Alkabani, and F. Koushanfar. “Realtime Emulations: Foundation and Applications.” *Design Automation Conference (DAC)*, 2010
29. [C] F. Koushanfar, **A. Mirhoseini**, and Y. Alkabani. “A Unified Submodular Framework for Multimodal IC Trojan Detection.” *Information Hiding Conference (IH)*, 2010

PEER REVIEWED
WORKSHOPS

1. **A. Mirhoseini**, B. Rouhani, E. Songhori, and F. Koushanfar. “ExtDict: Extensible Dictionaries for Data- and Platform-Aware Large-Scale Learning.” *IPDPS Workshop on Parallel and Distributed Computing for Large Scale Machine Learning and Big Data Analytics (ParLearning)*, 2017 ((**BEST PAPER AWARD**))
2. M. Samargh, **A. Mirhoseini**, and F. Koushanfar. “A Blocking Scheme for Deep Learning Acceleration.” *Bay Area Machine Learning Symposium (BayLearn)*, 2016
3. **A. Mirhoseini**, E. Dyer, R. Baraniuk, and F. Koushanfar. “Data- and Hardware-Aware Distributed Learning of Massive Non-Sparse Datasets.” *Annual Conference on Neural Information Processing Systems (NIPS) Workshop Series, BigNeuro*, 2015

RECENT INVITED
TALKS

1. Keynote: ML Workshop, Supercomputing Conference, 2018
2. Distinguished Speaker: Google AI/Tsinghua Symposium in China, 2018
3. UCSD ECE Seminar Series, 2018
4. International Supercomputing Conference, 2018
5. NIPS Deep Learning at Supercomputer Scale Workshop, 2017
6. Caltech, Dept. of Computing and Mathematical Sciences, 2016
7. Cornell, Dept. of Electrical and Computer Engineering, 2016

PATENTS

1. N. Shazeer, **A. Mirhoseini**, and K. Maziarz. ”Mixture of Experts Neural Networks.”
2. B. Rouhani, **A. Mirhoseini**, and F. Koushanfar. ”MobiDeep: Making Sense of Mobile Context by Deep Learning.” Under review for US patenting, Application No. 62294215, 2016
3. B. Rouhani, **A. Mirhoseini**, E. Songhori, and F. Koushanfar. ”Automated Real-Time Analysis of Streaming Big and Dense Data.” Under review for US patenting, Application No. 62329826, 2016
4. S. Nath, M. Goraczko, J. Liu, and **A. Mirhoseini**. “Optimizing Task Recommendations in Context-Aware Mobile Crowdsourcing.” Publication Number: 20150317582, 2015

COMPUTER SKILLS

- **Programming:** Python, MATLAB, C++, JavaScript
- **Parallel/Distributed Programming:** MPI, OpenMP, OpenCL, CUDA, MapReduce, familiar with GraphLab, Apache Spark
- **Operating Systems:** Linux, Windows, Mac OS

PROFESSIONAL
LEADERSHIP AND
SERVICES

- Co-Organizer, Machine Learning Day at International SuperComputing Conference, 2019
- Co-Organizer, ML for System Workshop, NeurIPS 2018
- Co-Organizer and Chair, IPDPS ParLearning Workshop, 2018
- Co-Organizer and Chair, ISCAS Scalable Deep Learning Workshop, 2017
- Executive Member and Web Chair, Design Automation Vision Challenge Sponsored by IEEE Computer-Aided Network DEsign (CANDE), DAC 2015
- Executive Committee Member and Event Organizer, **Women ExCEL** (Electrical and Computer Engineering Leaders), 2009-2015
- PC Member/Reviewed for:
 - Automation Conference (DAC), The Network and Distributed System Security Symposium (NDSS), IEEE International Conference on Computer Communications (INFOCOM), ACM Asia Conference on Computer and Communications Security (ASIACCS), International Conference on Applied Cryptography and Network Security (ACNS), NIPS Deep Learning Symposium, IEEE Transactions on Very Large Scale Integration (VLSI) Systems, IEEE Transactions on Forensics and Security (TIFS)