

Virginia Smith

Employment

- 2018– **Assistant Professor**, *Carnegie Mellon University*
- 2017–2018 **Postdoctoral Researcher**, *Stanford University*
- 2014 **Search Intern**, *Google*
- 2012–2017 **Research Assistant**, *UC Berkeley*

Education

- 2012–2017 **MS & PhD, Computer Science**, *University of California, Berkeley*
Advisors: Michael I. Jordan and David Culler
Thesis: *System-Aware Optimization for Machine Learning at Scale*
- 2008–2012 **BA, Mathematics & BA, Computer Science, Highest Distinction**, *University of Virginia*

Awards

- 2019 Carnegie Bosch Institute Research Award
- 2018 Google Faculty Research Award
- 2017 Outstanding GSI Award
- 2016 Rising Stars in EECS, Invited Participant
- 2015 MLconf Industry Impact Student Research Award
- 2014–2017 National Science Foundation Graduate Research Fellowship
- 2014 National Defense Science and Engineering Graduate Fellowship
- 2014 Tong Leong Lim Pre-Doctoral Prize
- 2014 Google Anita Borg Memorial Scholarship
- 2012–2014 Chancellor’s Fellowship
- 2012 College of Engineering Fellowship
- 2012 Department of Electrical Engineering and Computer Sciences Excellence Award
- 2012 CRA Outstanding Undergraduate Research Award, Honorable Mention
- 2012 Rader Undergraduate Research Award (UVA Top Undergraduate CS Research)
- 2008–2012 Echols Scholar (UVA Honors Program)
- 2008–2012 College Science Scholar (UVA Research Program)
- 2008 J. L. Wang Memorial Mathematics Scholarship

Publications

Preprints

M. Kuchnik, G. Amvrosiadis, and V. Smith, “Progressive compressed records: Taking a byte out of deep learning data,” <https://arxiv.org/abs/1911.00472>, 2019.

T. Li, A. Sahu, A. Talwalkar, and V. Smith, “Federated learning: Challenges, methods, and future directions,” <https://arxiv.org/abs/1908.07873>, 2019.

Refereed Conference or Journal

T. Li, M. Sanjabi, M. Zaheer, and V. Smith, “Fair resource allocation in federated learning,” in *International Conference on Learning Representations (ICLR)*, 2020.

- T. Li, A. Sahu, M. Sanjabi, M. Zaheer, A. Talwalkar, and V. Smith, “Federated optimization in heterogeneous networks,” in *Conference on Machine Learning and Systems (MLSys)*, 2020.
- T. Dao, A. Gu, A. Ratner, V. Smith, C. D. Sa, and C. Re, “A kernel theory of modern data augmentation,” in *International Conference on Machine Learning (ICML)*, 2019.
- M. Kuchnik and V. Smith, “Efficient augmentation via data subsampling,” in *International Conference on Learning Representations (ICLR)*, 2019.
- V. Smith, S. Forte, C. Ma, M. Takac, M. I. Jordan, and M. Jaggi, “CoCoA: A general framework for communication-efficient distributed optimization,” *Journal of Machine Learning Research*, 2018.
- V. Smith, C.-K. Chiang, M. Sanjabi, and A. Talwalkar, “Federated multi-task learning,” in *Neural Information Processing Systems (NIPS)*, 2017.
- C. Ma, J. Konecny, M. Jaggi, V. Smith, M. I. Jordan, P. Richtarik, and M. Takac, “Distributed optimization with arbitrary local solvers,” *Optimization Methods and Software*, 2017.
- V. Smith, M. Connor, and I. Stanton, “Going in-depth: Finding longform on the web,” in *Conference on Knowledge Discovery and Data Mining (KDD)*, 2015.
- C. Ma*, V. Smith*, M. Jaggi, M. I. Jordan, P. Richtarik, and M. Takac, “Adding vs. averaging in distributed primal-dual optimization,” in *International Conference on Machine Learning (ICML)*, 2015.
- M. Jaggi*, V. Smith*, M. Takac, J. Terhorst, S. Krishnan, T. Hofmann, and M. I. Jordan, “Communication-efficient distributed dual coordinate ascent,” in *Neural Information Processing Systems (NIPS)*, 2014.
- E. Sparks, A. Talwalkar, V. Smith, X. Pan, J. Gonzalez, T. Kraska, M. I. Jordan, and M. J. Franklin, “MLI: An API for user-friendly distributed machine learning,” in *International Conference on Data Mining (ICDM)*, 2013.
- J. Taneja, V. Smith, D. Culler, and C. Rosenberg, “A comparative study of high renewables penetration electricity grids,” in *IEEE International Conference on Smart Grid Communications*, 2013.
- A. Aswani, N. Master, J. Taneja, V. Smith, A. Krioukov, D. Culler, and C. Tomlin, “Identifying models of HVAC systems using semiparametric regression,” in *American Control Conference (ACC)*, 2012.
- V. Smith, T. Sookoor, and K. Whitehouse, “Modeling building thermal response to HVAC zoning,” *ACM SIGBED Review*, vol. 9, no. 3, 2012.

Invited Talks

- 2018 Machine Learning on Consumer Devices Workshop at NIPS
- 2018 Microsoft Research Faculty Summit
- 2017 ML Systems Workshop at NIPS
- 2017 Google Seattle
- 2017 Carnegie Mellon University
- 2017 Massachusetts Institute of Technology
- 2017 University of Washington
- 2017 University of California, Los Angeles
- 2017 Harvey Mudd College
- 2017 Microsoft Research, New England
- 2017 Microsoft Research, NYC
- 2017 Microsoft Research and MSR-NExT, Seattle
- 2017 SIAM Conference on Optimization (SIOPT)
- 2016 The Machine Learning Conference (MLconf)
- 2016 ML Systems Workshop at ICML
- 2015 The Machine Learning Conference (MLconf)
- 2015 Modeling and Optimization: Theory and Applications Conference (MOPTA)

- 2015 SF Bay Area Machine Learning Meetup
- 2014 Distributed Machine Learning and Matrix Computations Workshop at NIPS
- 2014 Bay Area Machine Learning Symposium (BayLearn)

Service and Activities

- 2019 **Co-Organizer**, *NeurIPS Workshop on Federated Learning for Data Privacy and Confidentiality*
- 2019 **Session Chair**, *Asilomar Session on Machine Learning and Optimization in Distributed Networks*
- 2018–2019 **Program Chair**, *SysML: Systems and Machine Learning Conference*
- 2017–2018 **Co-Organizer**, *SysML: Systems and Machine Learning Conference*
- 2015–2017 **Co-Founder and Organizer**, *Women in Technology Leadership Round Table (wit.berkeley.edu)*
- 2014–2015 **Co-President**, *UC Berkeley Women in Computer Science and Engineering (WICSE)*
- 2012–2014 **Officer**, *UC Berkeley Women in Computer Science and Engineering (WICSE)*

Reviewing

Neural Information Processing Systems (Best Reviewer Award 2017), International Conference on Machine Learning (Outstanding Reviewer 2018), Journal of Machine Learning Research, Foundations and Trends in Machine Learning, Conference on Systems and Machine Learning