

# Aditya Devarakonda |

3400 N. Charles Street – Baltimore, MD 21218  
☎ 732-306-9522 • ✉ devarakonda.aditya@gmail.com  
🌐 <https://www.aditya08.github.io>

## Education

---

### Ph.D. in Computer Science

*University of California, Berkeley*

2012–2018

Minor in Computational and Data Science and Engineering.

Thesis Committee: James Demmel, Michael W. Mahoney, and Adityanand Guntuboyina.

### M.S. in Computer Science

*University of California, Berkeley*

2012–2016

### B.S. in Computer Engineering

*Rutgers University, New Brunswick*

2008–2012

Double major in Computer Science.

## Interests

---

- Numerical Methods: Krylov subspace methods, First-order methods, Quasi-Newton methods.
- Parallel Algorithms: s-step methods, 2.5D numerical linear algebra.
- Machine Learning: convex optimization, deep learning.

## Work Experience

---

### JOHNS HOPKINS UNIVERSITY

*Assistant Research Scientist*

*Apr 2019 – Present*

Developing the theory and practice of efficient machine and deep learning algorithms. Applying deep learning methods to astrophysics.

### JOHNS HOPKINS UNIVERSITY APPLIED PHYSICS LABORATORY

*Data Scientist*

*Oct 2018 – Mar 2019*

Developing high performance code for data science problems faced by APL sponsors. Contributing developer to Socrates (distributed graph analytics) and pGalileo (A Parallel MPI-Based Generalized Low-Entropy Mixture Model). Technical lead for a high-performance computing systems modernization and procurement project.

### UNIVERSITY OF CALIFORNIA, BERKELEY

*Graduate Research Assistant*

*Aug 2012 – Aug 2018*

Designed a novel technique to avoid communication in parallel machine learning algorithms. Derived theoretical results to show the new algorithms are provably efficient and implemented the new algorithms in C++ and MPI to show speedups of up to  $6.1\times$  on a Cray XC30 supercomputer.

### NVIDIA CORPORATION

*Research Intern*

*Aug 2017 – Nov 2017*

Derived an adaptive batch size technique which progressively increases the batch size during the training of deep neural networks. Implemented the technique in the PyTorch deep learning framework and experimentally showed speedups of up to  $6.25\times$  when using 4 NVIDIA P100 and V100 GPUs.

### RUTGERS UNIVERSITY, NEW BRUNSWICK

*Undergraduate Research Assistant*

*Sept 2010 – May 2012*

Developed a scheduling and communication framework for replica exchange molecular dynamics to be deployed on a heterogeneous environment comprised of Cloud, Supercomputing, and GRID computing resource.

## UNIVERSITY OF ILLINOIS, URBANA-CHAMPAIGN

Research Assistant

May 2011 – Aug 2011

Developed a parallel Quicksort algorithm and topology-aware dynamic load balancing techniques for ChaNGa, an n-body cosmology application.

## Teaching Experience

---

### JOHNS HOPKINS UNIVERSITY CENTER FOR TALENTED YOUTH INTERNATIONAL (SOUTH KOREA)

Workshop Presenter - Introduction to Deep Learning

January 2020

Developed interactive workshop material using Jupyter notebooks intended to be accessible for students without any computer science or machine learning background.

### JOHNS HOPKINS UNIVERSITY

Guest Lecturer - AS117.205 Beautiful Data: Introduction to Practical Data Science

Spring 2019 & Spring 2020

Gave two introductory lectures on supervised and unsupervised learning which covered the theory and practice. Provided demonstrations of supervised and unsupervised learning in action.

### UNIVERSITY OF CALIFORNIA, BERKELEY

Graduate Student Instructor - CS267 Applications of Parallel Computers Spring 2014, Summer 2016 & Spring 2018

Graded projects, assignments, and developed autograders. Developed online courseware in Summer 2016 which was used to teach for-credit classes at several universities in the country through XSEDE. Served as the head graduate student instructor (GSI) in 2018 and worked with the instructors and other GSIs to offer three versions of the course concurrently. Mentored two graduate students in Spring 2018 on implementing communication-avoiding block coordinate descent on Spark.

### UNIVERSITY OF CALIFORNIA, BERKELEY

Guest Lecturer - CS267 Applications of Parallel Computers

Spring 2017

Gave one introductory lecture on communication-avoiding algorithms for machine learning with a focus on the block coordinate descent method.

### PARALLEL COMPUTING LAB (PARLAB) - UNIVERSITY OF CALIFORNIA, BERKELEY

Teaching Assistant - Short Course on Parallel Computing

Summer 2014

Co-led hands-on parallel programming sessions, moderated discussions, and presented tutorials on parallel programming models (OpenMP, MPI, CUDA).

### JOHNS HOPKINS UNIVERSITY CENTER FOR TALENTED YOUTH

Teaching Assistant - Intro. to Electrical Engineering

Summer 2010 & Summer 2014

Led lab and study sessions, graded assignments, quizzes and reports. Designed and taught lessons on C++ programming and Arduino programming for audience of middle school students.

### RUTGERS UNIVERSITY, NEW BRUNSWICK

Teaching Assistant - CS111 Intro. to Computer Science

Spring 2010

Led several review sections and designed lessons to prepare for exams and assignments.

## Mentorship

---

- Alex Michon. CS267 final project on implementing and benchmarking communication-avoiding machine learning algorithms on Apache Spark.
- Jieming Wei. CS267 final project on implementing and benchmarking communication-avoiding machine learning algorithms on Apache Spark.

## Thesis

---

- A. Devarakonda, Avoiding Communication in First Order Methods for Optimization, Ph.D. Thesis, University of California, Berkeley, 2018.
- A. Devarakonda, Exploring Asynchronous Application Formulations for Heterogeneous Infrastructure, Undergraduate Thesis, Rutgers University, New Brunswick, 2012.

## Manuscripts in Progress

---

- Avoiding Communication in Kernel Methods: Regression and Classification.

## Publications and Technical Reports

---

- A. Devarakonda, and J. Demmel, Avoiding Communication in Logistic Regression, International Conference on High Performance Computing, Data, and Analytics (HiPC), 2020. **(To Appear)**
- A. Devarakonda, K. Fountoulakis, J. Demmel and M. W. Mahoney, Avoiding communication in primal and dual block coordinate descent methods, SIAM Journal of Scientific Computing (SISC), 41(1), pp. C1-C27, 2019.
- A. Devarakonda, K. Fountoulakis, J. Demmel and M. W. Mahoney, Avoiding Synchronization in First-Order Methods for Sparse Convex Optimization, 2018 In Proceedings of the IEEE International Parallel and Distributed Processing Symposium (IPDPS 2018), Vancouver, British Columbia, Canada, 2018, pp. 409-418.
- S. Soori, A. Devarakonda, Z. Blanco, J. Demmel, M. Gurbuzbalaban, and M. M. Dehnavi. Reducing Communication in Proximal Newton Methods for Sparse Least Squares Problems. In Proceedings of the 47th ACM International Conference on Parallel Processing (ICPP 2018), New York, NY, USA.
- A. Devarakonda, M. Naumov, and M. Garland, AdaBatch: Adaptive Batch Sizes for Training Deep Neural Networks, arXiv cs.LG 1712.02029, 2017.
- S. Soori, A. Devarakonda, J. Demmel, M. Gurbuzbalaban, M. M. Dehnavi, Avoiding communication in proximal methods for convex optimization problems, arXiv cs.DC 1710.08883, 2017
- A. Gittens, A. Devarakonda, E. Racah, M. Ringenburt, L. Gerhardt, J. Kottalam, J. Liu, K. Maschhoff, S. Canon, J. Chhugani, P. Sharma, J. Yang, J. Demmel, J. Harrell, V. Krishnamurthy, M. W. Mahoney, Prabhat, Matrix factorizations at scale: A comparison of scientific data analytics in spark and C+MPI using three case studies, 2016 IEEE International Conference on Big Data (Big Data 2016), Washington, DC, 2016, pp. 204-213.
- R. Carbutescu, A. Devarakonda, J. Demmel, S. Gordon, J. Alameda, and S. Mehringer. 2014. Architecting an autograder for parallel code. In Proceedings of the 2014 Annual Conference on Extreme Science and Engineering Discovery Environment (XSEDE 2014). ACM, New York, NY, USA.
- M. Parashar, M. AbdelBaky, I. Rodero and A. Devarakonda, Cloud Paradigms and Practices for Computational and Data-Enabled Science and Engineering, in Computing in Science & Engineering, vol. 15, no. 4, pp. 10-18, 2013.
- D. Villegas, N. Bobroff, I. Rodero, J. Delgado, Y. Liu, A. Devarakonda, L. Fong, S. M. Sadjadi, M. Parashar, Cloud federation in a layered service model, Journal of Computer and System Sciences, Volume 78, Issue 5, pp. 1330-1344, 2012.
- A. Devarakonda, I. Rodero, M. AbdelBaky, B. Claus, M. Parashar, Scalable Asynchronous Molecular Dynamics on Federated Cyberinfrastructure, Poster in XSEDE Conference, July 2012.

## Presentations

---

- Avoiding Communication in First-Order Methods for Optimization, Invited Talk, Reservoir Labs, January 2021.
- Avoiding Communication in Logistic Regression, International Conference on High Performance Computing, Data, and Analytics, December 2020.
- Unsupervised Learning, Guest Lecturer, AS.171.205 Beautiful Data: Introduction to Practical Data Science, Johns Hopkins University, April 2019.
- An Introduction to Deep Learning, Guest Lecturer, AS.171.205 Beautiful Data: Introduction to Practical Data Science, Johns Hopkins University, April 2019.
- Avoiding Communication in Machine Learning, CS Seminar, Lawrence Berkeley National Lab, Berkeley, CA, January 2019.

- Avoiding Synchronization in First-Order Methods for Sparse Convex Optimization, International Parallel and Distributed Processing Symposium, May 2018.
- Avoiding Communication in First-Order Methods for Optimization, Dissertation Talk, Berkeley, CA, May 2018.
- s-step Methods in Machine Learning, SIAM Parallel Processing, Tokyo, Japan, March 2018.
- Avoiding Synchronization in Sparse Convex Optimization, UC Berkeley Scientific Computing and Matrix Computations Seminar, Berkeley, CA, February 2018.
- Avoiding Communication in Machine Learning, Johns Hopkins University Applied Physics Lab, Laurel, MD, January 2018.
- Communication-Avoiding Krylov Subspace Methods, NASA Langley Research Center, Hampton, VA, December 2017.
- Communication-Avoiding Algorithms, NASA Langley Research Center, Hampton, VA, December 2017.
- AdaBatch: Adaptive Batch Sizes for Training Deep Neural Networks, NVIDIA, Santa Clara, CA, November 2017.
- Communication-Avoiding Methods for Regularized Least-Squares, SIAM Annual, Pittsburgh, PA, July 2017.
- Communication-Avoiding Machine Learning, Householder Symposium XX, Blacksburg, VA, June 2017.
- Communication-Avoiding Machine Learning, ASPIRE Summer Retreat, Santa Cruz, CA, June 2017.
- Communication Avoiding Primal and Dual Block Coordinate Descent Methods, SIAM Optimization, Vancouver, Canada, May, 2017.
- A Generalized Framework for Communication-Avoiding Regularized Least Squares, Berkeley Statistics Annual Research Symposium (BSTARS), Berkeley, CA, April, 2017.
- Communication-Avoiding Machine Learning: Block Coordinate Descent, Guest Lecturer, CS267 Applications of Parallel Computers, University of California, Berkeley, March, 2017.
- Matrix factorizations at scale: A comparison of scientific data analytics in Spark and C+ MPI using three case studies, IEEE Big Data, Washington D.C., December, 2016.
- Avoiding Communication in Machine Learning, Berkeley Statistics Annual Research Symposium (BSTARS), Berkeley, CA, 2016.
- Communication-Avoiding Coordinate Descent Methods for Linear Systems, UC Berkeley Scientific Computing and Matrix Computations Seminar, Berkeley, CA, 2016
- Scalable Asynchronous Molecular Dynamics on Federated Cyberinfrastructure, SCALE Challenge 2012, Ottawa, Canada, May, 2012.
- Parallel Quicksort, Passionate on Parallel Research Experience for Undergraduates, Urbana, IL, August, 2011.
- Molecular Dynamics on Federated Cyberinfrastructure, Demo for Life Sciences and Pharmaceuticals Industrial Advisory Board, New Brunswick, NJ 2011.

## Computational Resource Grants

---

**Cori (Cray XC40): 148,088 hours**  
NERSC

**Comet: 90,000 hours, Stampede2: 10,000 hours**  
XSEDE

## Professional Service

---

**UC Berkeley Graduate Admissions 2015:** Student Reviewer

**UC Berkeley Faculty Search 2017:** Student Reviewer

**TOPC 2018:** Reviewer

## Outreach

---

### **EECS PEERS**

*Computer Science Coordinator*

### **BIAS BUSTERS**

*Workshop Presenter*

### **UC BERKELEY GRADUATE ASSEMBLY**

*EECS Delegate*

### **INFORMATION RISK GOVERNANCE COMMITTEE**

*Committee Member*

### **CAMPUS INFORMATION SECURITY AND PRIVACY**

*Committee Member*

## Selected Awards

---

- NSF Graduate Research Fellowship
- UC Berkeley EECS Department Fellowship
- Rutgers School of Engineering James J. Slade Scholar
- James L. Potter Award in Research
- Kuhl Memorial Engineering Fellowship
- Charles V. Longo Memorial Scholarship
- Rutgers Academic Excellence Award
- Hannah Sands Endowed Scholarship in Engineering
- Rutgers School of Engineering Dean's List
- Edward J. Bloustein Distinguished Scholars Award
- Rutgers Scarlet Scholarship

## Travel Awards

---

- IPDPS 2018 Conference Travel Award
- Householder XX Symposium Travel Award
- SIAM Parallel Processing 2018 Student Travel Award
- UC Berkeley Conference Travel Grant (2016, 2017)