

REZA ZADEH

rezab@stanford.edu

reza-zadeh.com

650.898.3193

SUMMARY

Founder CEO at Matroid, Adjunct Professor at Stanford University
Served on the Technical Advisory board of Microsoft and Databricks
Ph.D. from Stanford. Masters from Carnegie Mellon. Bachelors from Waterloo.
Advanced research and engineering positions at Google, Twitter, IBM
Focused on Machine Learning, Distributed Computing, and Discrete Applied Mathematics
Contributor to Apache Spark and TensorFlow, initial creator of the Linear Algebra Package in Spark
Wrote an O'Reilly Media book on TensorFlow, used at Matroid for Computer Vision
Created and taught courses on distributed algorithms and discrete mathematics at Stanford
Mentored undergraduate, masters, and Ph.D. students on research
Sifted through petabytes of data by way of Google MapReduce and Apache Hadoop, Pig, Spark
Published and refereed for Machine Learning conferences and journals
Organized international conferences and workshops on Machine Learning
Interactive resume at <http://reza-zadeh.com>.

EDUCATION

Ph.D. in COMPUTATIONAL MATHEMATICS

Stanford University

Advisor: Gunnar Carlsson

April 2014

Thesis: Large Scale Graph Completion.

Master of COMPUTER SCIENCE

Carnegie Mellon University

August 2010

Fully funded research masters.

Bachelor of MATHEMATICS, Honors COMPUTER SCIENCE, SOFTWARE ENGINEERING

University of Waterloo

August 2008

Active in undergraduate research and programming competitions.

EXPERIENCE

Matroid

Palo Alto, California

January 2016 - Current

Chief Executive Officer

- Making Computer Vision Simple. \$15.5 million raised from Intel, New Enterprise Associates (NEA), and others.

Microsoft

Redmond, Washington

December 2015 - November 2016

Technical Advisor

- Advising the leaders of the Applications and Services group as a member of Microsoft's Technical Advisory Board on Machine Learning and Recommender Systems

Databricks

Berkeley, California

November 2013 - November 2016

Technical Advisor

- Advising the founders on all areas of distributed machine learning and numerical linear algebra
- Building Spark's machine learning library: MLlib

Stanford University

Stanford, California

July 2014 - Current

- Various roles: currently Adjunct Professor

- Conducting research on distributed matrix computations and machine learning
- Teaching several classes on discrete mathematics and distributed algorithms
- Co-advising graduate students (masters and Ph.D.)
- Organizing international workshops and conferences on Machine Learning
- Leading C2 and SMACC Consulting: stanford.edu/~rezab/smacc

Twitter

San Francisco, California

May 2011 - April 2013

Senior Data Scientist

- Responsible for the Machine Learning behind http://twitter.com/who_to_follow
- Pushed the DIMSUM and DISCO algorithms into production
- Open-sourced DIMSUM and DISCO algorithms via Spark
- Produced Web Graph on Twitter's slice of the web
- Wrote and ran PageRank, personalized PageRank and its flavors on web graph, pushed to production
- Modeled the timing of notifications using the Secretary Problem
- Mentored intern (Sebastian Schelter). Project: a distributed parameter-server for matrix completion

Google

Headquarters, Mountain View, California

January 2006 - September 2007

Research

- Wrote and ran some of the world's largest MapReduce jobs
- Improved machine translation quality using several bridge languages
- Reduced trillion-word Language Model (LM) size for SMT by 20%, without affecting quality
- Sifted through terabytes of data by way of: MapReduce (Hadoop), GFS, SSTables, BigTable

Morgan Stanley

Headquarters, New York, NY

January 2008 - May 2008

Financial Engineering

- Wrote algorithms to price single-name Credit Default Swaps, back out default probabilities, then price more exotic instruments

IBM (International Business Machines)

Markham Software Lab, Ontario, Canada

May 2005 - August 2005

Software Engineering

- Built plugins for Eclipse and Websphere Development Studio
- Designed an automated update system allowing WebSphere Development Studio Client 7.0 to keep AS400 machines up to date

TEACHING EXPERIENCE

PRIMARY INSTRUCTOR, FULL QUARTER CLASS

- CME 305: Discrete Mathematics and Algorithms (MS&E 316) (38 students)
Responsible for the associated Ph.D. qualifying exam.
Class evaluation: 4.6 out of 5.0
- MS&E 317: Algorithms for Modern Data Models (CS 263) (20 students)
Class evaluation: 4.27 out of 5.0
- CME 323: Distributed Algorithms and Optimization (55 students)
- CME 400/291: Ph.D./Masters Directed Research (3 students)
- Many more classes and specific offerings listed on my website <http://stanford.edu/~rezab>

PRIMARY INSTRUCTOR, SHORT COURSE

- Advanced Data Science, Spark Summit (240 students)
Evaluation: 8.5 out of 10

- Stanford Spark class (500 students)
Class evaluation: 6.83 out of 7.0
- Stanford Spark workshop (250 students)
Class evaluation: 6.4 out of 7.0
- University of Maryland Spark Class (140 students)
Class evaluation: 6.4 out of 7.0
- Linear Algebra Refresher Class (20 students)
- Discrete Mathematics Refresher Class (20 students)

TEACHING ASSISTANT

- Linear Algebra (CME 200)
- Discrete Mathematics and Algorithms (CME 305)

ADVISING EXPERIENCE

- Mentoring Sebastian Schelter (Twitter intern) on building a distributed parameter server for matrix completion.
I am on Sebastian's Ph.D. Thesis Committee at TU Berlin.
- Advising Vishakh Hegde (CME 291)
Vishakh and I obtained leading results on the Princeton ModelNet competition.
- Advising Arun Jambulapati (Stanford M.S. student) on the planted clique problem.
Arun is taking M.S. Directed Research (CME 291) with me, and TAed CME 305 for me.
- Advising Burak Yavuz (Stanford masters student) on building block matrix multiplication on Spark.
Burak took Masters Directed Research (CME 291) with me, and was my student in CS 263. Our paper won a best paper award at KDD 2016.
- Helping many students and researchers via one-on-one discussions of their research problems while leading the Stanford SMACC & C2 consulting service: <http://stanford.edu/~rezab/smacc>

CONFERENCES ORGANIZED

1. Scaled Machine Learning Conference, Stanford 2016
2. Distributed Machine Learning and Matrix Computations Workshop, NIPS 2014
Organized with Ion Stoica and Ameet Talwalkar
3. Modern Massive Data Sets 2014 (MMDS 2014), at UC Berkeley
Organized with Michael Mahoney, Alexander Shkolnik, Petros Drineas, Fernando Perez
4. Large-scale Matrix Analysis and Inference Workshop, NIPS 2013
Organized with Gunnar Carlsson, Wouter M. Koolen, Michael Mahoney, Manfred Warmuth
5. Clustering Theory Workshop, NIPS 2009
Organized with Shai Ben-David, Ulrike von Luxburg, Avrim Blum, Isabelle Guyon, Robert C. Williamson, Margareta Ackerman

PEER-REVIEW SERVICE

Journal of Machine Learning Research
Neural Information Processing Systems (NIPS)
Knowledge Discovery and Data Mining (KDD)
Pattern Recognition Journal
European Conference on Machine Learning and PKDD
Stanford Aeronautics & Astronautics Department Faculty Search Student Committee

HONORS & AWARDS

KDD 2016 Best Paper Award
Gene Golub Outstanding Thesis Award, for best thesis in department
ICME Student Leadership award
Canada Graduate Scholarship Doctoral
Alexander Graham Bell Graduate Scholarship
Queen Elizabeth II Aiming for the Top
Google Code Jam 2006 Onsite Finalist, New York, NY
Waterloo Dean's Honor List

SELECTED TALKS

1. MLlib and Distributing the Singular Value Decomposition, Stanford University
2. Dimension Independent Matrix Square, MMDS 2014, UC Berkeley
3. The Libraries of Spark, Keynote at Data Science Bootcamp
4. Trends in Big Data, Stanford University
5. Distributed Computing with Spark, University of Maryland
6. MLlib and All-Pairs Similarity, University of Maryland
7. Distributed Computing with Spark, eBay, Bay Area ACM
8. Towards a Principled Theory of Clustering, Carnegie Mellon University
9. Dimension Independent Matrix Square, Topology Day, Ayasdi
10. Distributed Computing with Spark, Twitter
11. Introduction to all-pairs similarity and SVD in Spark MLlib, Codeneuro San Francisco
12. Distributed Computing with Spark, ICME Xtend, Stanford University
13. Apache Spark in Four Parts, Raytheon
14. Learning Who to Follow, Twitter
15. Data-flow languages and Open-Source software, Infosys
16. A Uniqueness Theorem for Clustering, NIPS 2009 Conference
17. Supervised Clustering, NIPS 2010 Conference
18. Machine Learning on Spark, LinkedIn
19. Dimension Independent Matrix Square using MapReduce, NIPS 2013 Workshop
20. Machine Learning on Spark, Big Data Montreal
21. Machine Learning on Spark, Toronto Hadoop Users Group
22. Dimension Independent Matrix Square, McGill University
23. Spark Camp: An Introduction to Apache Spark with Hands-on Tutorials, Strata 2015
24. Introduction to Data Science with Spark, Spark Summit East 2015
25. Panel on Future of Deep Learning at Innovation Endeavors
26. Host of Spark Session, Strata 2016, San Jose
27. Matrix Computations and Optimization in Apache Spark, KDD 2016 plenary session

PUBLICATIONS

CONFERENCE PUBLICATIONS

1. Reza Zadeh and Shai Ben-David. A uniqueness theorem for clustering. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 639–646. AUAI Press, 2009
2. Pankaj Gupta, Ashish Goel, Jimmy Lin, Aneesh Sharma, Dong Wang, and Reza Zadeh. Wtf: The who to follow service at twitter. In *Proceedings of the 22nd international conference on World Wide Web*, pages 505–514. International World Wide Web Conferences Steering Committee, 2013
3. Reza Zadeh and Gunnar Carlsson. Dimension independent matrix square using mapreduce. *Poster at Symposium on Theory of Computing (STOC 2013)*, 2013
4. Lynne Burks, Mahalia Miller, and Reza Zadeh. Rapid estimate of ground shaking intensity by combining simple earthquake characteristics with tweets. *National Conference on Earthquake Engineering*, 2014
5. Pranjal Awasthi and Reza Zadeh. Supervised clustering. In *Advances in Neural Information Processing Systems*, pages 91–99, 2010
6. Aruna D Balakrishnan, Sara Kiesler, Jonathon Cummings, and Reza Zadeh. Research team integration: What it is and why it matters. In *Proceedings of the ACM 2011 conference on Computer supported cooperative work*, pages 523–532. ACM, 2011
7. Reza Zadeh, Aruna Balakrishnan, Sara Kiesler, and Jonathon Cummings. What’s in a move?: normal disruption and a design challenge. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2897–2906. ACM, 2011
8. Reza Zadeh, Ashish Goel, Kamesh Munagala, and Aneesh Sharma. On the precision of social and information networks. In *Proceedings of the ACM conference on Online social networks*, pages 63–74. ACM, 2013
9. Reza Bosagh Zadeh, Xiangrui Meng, Burak Yavuz, Aaron Staple, Li Pu, Shivaram Venkataraman, Evan Sparks, Alexander Ulanov, and Matei Zaharia. Matrix computations and optimization in apache spark. *Knowledge Discovery and Data Mining (KDD)*, 2016

JOURNAL PUBLICATIONS

1. Reza Zadeh and Ashish Goel. Dimension independent similarity computation. *The Journal of Machine Learning Research*, 14(1):1605–1626, 2013
2. Trevor Hastie, Rahul Mazumder, Jason D. Lee, and Reza Zadeh. Matrix completion and low-rank svd via fast alternating least squares. *The Journal of Machine Learning Research*, to appear, 2015
3. Jonathon N Cummings, Sara Kiesler, Reza Zadeh, and Aruna D Balakrishnan. Group heterogeneity increases the risks of large group size a longitudinal study of productivity in research groups. *Psychological science*, 24(6):880–890, 2013
4. Reza Zadeh. Machine learning on big data: How apache spark can help. *Biomedical Computation Review*, 2015
5. Xiangrui Meng, Joseph Bradley, Burak Yavuz, Evan Sparks, Shivaram Venkataraman, Davies Liu, Jeremy Freeman, DB Tsai, Manish Amde, Sean Owen, Doris Xin, Reynold Xin, Michael J. Franklin, Reza Zadeh, Matei Zaharia, and Ameet Talwalkar. Millib: Machine learning in apache spark. *Journal of Machine Learning Research - Open Source Software*, 2015
6. Madeleine Udell, Corinne Horn, Reza Zadeh, and Stephen Boyd. Generalized low rank models. *Foundations and Trends™ in Machine Learning*, 2016

INDUSTRY POSTS AND PUBLICATIONS

1. Reza Bosagh Zadeh. On the evolution of machine learning: from linear models to neural networks. *O’Reilly Radar*, 2015

2. Li Pu and Reza Zadeh. Distributing the singular value decomposition with spark. *Databricks Engineering Blog*, 2014
3. Reza Zadeh. Efficient similarity algorithm now available in spark, thanks to twitter. *Databricks Engineering Blog*, 2014
Covered by: Gigaom
4. Reza Zadeh. Using twitter to measure earthquake impact in almost real time. *Twitter Engineering Blog*, 2014
Covered by: TechCrunch, Mashable, Engadget, VentureBeat, Scientific American, and many other news outlets
5. Reza Zadeh. All-pairs similarity via dimsum. *Twitter Engineering Blog*, 2014
Covered by: Gigaom
6. Reza Zadeh. Dimension independent similarity computation (disco). *Twitter Engineering Blog*, 2012
Translated into Chinese by Xu Wenhao

WORKSHOP PUBLICATIONS

1. Sebastian Schelter, Venu Satuluri, and Reza Zadeh. Factorbird - a parameter server approach to distributed matrix factorization. *NIPS 2014 Workshop on Distributed Machine Learning and Matrix Computations*, 2014
2. Madeleine Udell, Corinne Horn, Reza Zadeh, and Stephen Boyd. Generalized low rank models. *NIPS 2014 Workshop on Distributed Matrix Computations*, 2014
3. Reza Zadeh and Gunnar Carlsson. Dimension independent matrix square using mapreduce. *NIPS 2013 Workshop on Large-scale matrix analysis*, 2013

THESIS

1. Reza Zadeh. *Large Scale Graph Completion*. PhD thesis, Stanford University, 2014

BOOK

1. Bharath Ramsundar and Reza Zadeh. *TensorFlow for Deep Learning: from Linear Regression to Reinforcement Learning*. O'Reilly Media, 2018

MEDIA COVERAGE

ENGLISH

1. Derrick Harris. Twitter open sourced a recommendation algorithm for massive datasets. *Gigaom*, 2014
2. Jon Fingas. Stanford turns to twitter to track earthquakes. *Engadget*, 2014
3. Catherine Shu. Tweets can guide emergency responders almost immediately after an earthquake. *TechCrunch*, 2014
4. Kurt Wagner. Can studying tweets lead to faster earthquake recovery? *Mashable*, 2014
5. M.G. Siegler. Twitter's social graph is about to get pumped up. "who to follow" is social steroids. *TechCrunch*, 2011
6. M.G. Siegler. Like facebook, twitter starts using algorithms to bulk up social graph. *TechCrunch*, 2011
7. M.G. Siegler. With tweaks, twitter "who to follow" gets stronger, more agile. *TechCrunch*, 2011
8. Kia Kokalitcheva. Stanford scientists use twitter to make more accurate earthquake maps. *VentureBeat*, 2014
9. Nicole Bogart. Stanford researchers use twitter data to create more accurate earthquake maps. *Global News*, 2014

10. Becky Oskin. #earthquake! tweets beat official quake alerts. *Live Science*, 2014
11. Ariana Whitmore. Stanford attempts to track the tremors of earthquake via twitter. *The Next Digit*, 2014

OTHER LANGUAGES

1. Franziska Weiss. Stanford erforscht erdbeben-analyse via twitter. *Engadget DE*, 2014
2. Nicola Di Turi. Twitter non prevede i terremoti ma. *Corriere*, 2014
3. M. Jara. Stanford utilizara a twitter para seguir terremotos. *Latercera*, 2014
4. Flvio Macedo. Twitter utilizado para fazer mapas mais precisos de terremoto. *Zuti*, 2014
5. B.V. Twitter i stanford imaju plan. *dnevik.hr*, 2014
6. ASTICun. Shakemaps : le programme des chercheurs de stanford utilise twitter. *Developpez*, 2014