

David P. Woodruff

Curriculum Vitae

Biographical

Computer Science Department
Gates-Hillman Complex
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213

Citizenship: United States

Email: dwoodruf@cs.cmu.edu

Home Page: <http://www.cs.cmu.edu/~dwoodruf/>

Research Interests

Compressed Sensing, Data Stream Algorithms and Lower Bounds, Distributed Computation, Machine Learning, Numerical Linear Algebra, Sketching

Education

All degrees received from Massachusetts Institute of Technology, Cambridge, MA

Ph.D. in Computer Science. September 2007

Research Advisor: Piotr Indyk

Thesis Title: Efficient and Private Distance Approximation in the Communication and Streaming Models.

Master of Engineering in Computer Science and Electrical Engineering, May 2002

Research Advisor: Ron Rivest

Thesis Title: Cryptography in an Unbounded Computational Model

Bachelor of Science in Computer Science and Electrical Engineering May 2002

Bachelor of Pure Mathematics May 2002

Professional Experience

August 2017 - present, Carnegie Mellon University
Computer Science Department
Associate Professor (with tenure)

June 2018 – December 2018, Google, Mountain View
Research Division
Visiting Faculty Program

August 2007 - August 2017, IBM Almaden Research Center
Principles and Methodologies Group
Research Scientist

Aug 2005 - Aug 2006, Tsinghua University
Institute for Theoretical Computer Science
Visiting Scholar. Host: Andrew Yao

Jun - Jul 2005, DoCoMo Research Labs
Cryptography and Security Group
Research Intern. Mentors: Craig Gentry and Zulfikar Ramzan

May - Aug 2003, Palo Alto Research Center
Computer Science Lab
Research Intern. Mentor: Jessica Staddon

Jun - Aug 1999, January 2000, June - Aug. 2000, Hughes Research Labs
Signal Processing Group
Research Intern. Host: Peter Capofreddi

Awards

- Simons Investigator Award, 2020
- PODS 2020 Best Paper Award for the paper “A Framework for Adversarially Robust Streaming Algorithms (coauthors Omri Ben-Eliezer, Rajesh Jayaram, and Eylon Yogev)
- 2014 EATCS Presburger Award
<http://www.eatcs.org/index.php/component/content/article/1-news/1866-presburger-award-2014/>
- STOC 2013 Best Paper Award for the paper “Low Rank Approximation and Regression in Input Sparsity Time” (coauthor Ken Clarkson)
Paper invited to the Journal of the ACM
- PODS 2010 Best Paper Award for the paper “An Optimal Algorithm for the Distinct

Elements Problem” (coauthors Daniel M. Kane, Jelani Nelson)
Paper invited to the Journal of the ACM

- Elected into the IBM Academy of Technology in 2014
http://en.wikipedia.org/wiki/IBM_Academy_of_Technology
- IBM Master Inventor Award in 2012, an award given to an employee based on a significant contribution to the IBM Patent Portfolio, see:
http://en.wikipedia.org/wiki/IBM_Master_Inventor
Multiple IBM Achievement Plateau Awards for patents filed
- “Outstanding” IBM Research Division Accomplishment awarded in 2019 for earlier research on implementing fast sketching techniques in the Skylark package, as part of the XDATA project
- “Outstanding” IBM Research Division Accomplishment for 2012 for research on “Processing Large Datasets: Algorithms and Computational Limits” (one award given to all of IBM Almaden Computer Science in 2012)
- IBM Research Division Accomplishment for 2014 for research on “Low Complexity Kernels for Cognitive Computing”
- IBM Pat Goldberg Award for one of the four Best IBM Research Papers in Computer Science, Electrical Engineering and Math published in 2013 for the paper “Low Rank Approximation and Regression in Input Sparsity Time” (coauthor Ken Clarkson)
- IBM Pat Goldberg Award for one of the five Best IBM Research Papers in Computer Science, Electrical Engineering and Math published in 2012 for the paper “Sublinear Optimization for Machine Learning” (coauthors Ken Clarkson, Elad Hazan)
- IBM Pat Goldberg Award for one of the four Best IBM Research Papers in Computer Science, Electrical Engineering and Math published in 2010 for the paper “An Optimal Algorithm for the Distinct Elements Problem” (coauthors Daniel M. Kane, Jelani Nelson)
- DOD/NDSEG Graduate Fellowship 2003-2007
- Akamai Presidential Fellow 2002-2003

Professional Activities

- Program Chair for 2018 Program on Foundations of Data Science at the Simons Institute at UC Berkeley
- Program Committee Memberships:

- 2020: SODA, ITCS, ICML area chair, IJCAI senior PC, ESA, NeurIPS area chair
 - 2019: COLT, NeurIPS area chair, ICML, RANDOM, ISIT, IJCAI senior PC
 - 2018: STOC, SODA, ICML, ITCS, AAAI senior PC, ISIT, SOSA, NeurIPS, ICDT
 - 2017: RANDOM, KDD senior PC, SPAA, ICML, AISTATS
 - 2016: WSDM, NIPS, KDD, ICML, PODS
 - 2015: MASSIVE, RANDOM, ICDT, SIGMOD/PODS Ph.D. Symposium
 - 2014: FOCS, ESA, SPIRE
 - 2013: PODS, RANDOM, SPIRE, TAMC
 - 2012: FOCS, SODA, PODS, FAW
 - 2008: RANDOM
- Editorial Boards
 - Transactions on Algorithms (2015 – present)
 - Journal of Computer System Sciences (2016-present)
 - Data Science Book Series (2018-present)
 - Stochastic Models (2019-present)
- Guest Editorships
 - Co-editor for Transactions on Algorithms special issue for SODA, 2018
 - Co-editor for Algorithmica issue on Information Complexity and Applications
 - Co-editor for Transactions on Algorithms special issue for SODA, 2012
- Workshops Organized
 - Co-Organizer for session in the 2020 Information Theory and Applications (ITA) workshop on Low Rank Approximation
 - Co-organizer of 2020, Theoretical Computer Science Visioning Workshop (upcoming)
 - Co-organizer of 2019 Foundations of Data Science Reunion Workshop at Simons
 - Co-organizer of DIMACS and Simons Special Focus Workshop on Bridging Discrete and Continuous Optimization for 2019
 - Co-organizer of Data Stream Algorithms Workshop at Institute for Mathematical Sciences at Nanyang Technological Institute, Singapore, 2020 (upcoming)
 - Co-organizer for 2018 Foundations of Data Science Bootcamp at Simons
 - Organizer for 2018 Data Science Finale Workshop at Simons
 - Co-Organizer for FOCS 2017 Workshop on Linear Sketching
 - Co-Organizer for Shonan 2017 Workshop on Processing Big Data Streams
 - Co-Organizer for Dagstuhl 2017 Workshop on Theory and Applications of Hashing
 - Co-Organizer for Shonan 2016 Workshop on Recent Advances in Randomized Numerical Linear Algebra
 - Co-Organizer for 2015 Program in Computer Science for the Institut Henri Poincare in Paris for the “Inference Workshop” on streaming, sketching, sublinear algorithms, and compressed sensing
 - Co-Organizer of Workshop on Algorithms for Data Streams, Dortmund, Germany, 2012

- Co-Organizer for session in the 2016 Information Theory and Applications (ITA) workshop on Heavy Hitters
- Co-Organizer for session in the 2015 Information Theory and Applications (ITA) workshop on Sketching for Numerical Linear Algebra
- National Science Foundation (NSF) Panel Service for AF – Algorithms and Foundations
- Member of the SIGACT Committee for the Advancement of Theoretical Computer Science (CATCS), 2017-present
- IBM Raviv Fellowship Selection Committee Co-Chair, 2012 and 2013
- Simons Institute Long-Term Fellow in:
 - Theoretical Foundations of Big Data Analysis Program, Fall 2013
 - Theoretical Foundations of Machine Learning, Spring 2017
 - Foundations of Data Science (chair), Fall 2018
- CMU graduate student admissions committee
 - 2018
 - 2019
- CMU faculty review and promotion committee reader
- Extensive journal and conference reviewing. Selected for top 5% of ICML reviewers in 2019.
- External reviewer for university hiring
- IBM committee service, e.g., computer science workshop and social planning. CMU community activities, e.g., SCS4All Faculty Ask Me Anything Panel

Summer Interns Mentored

- Jelani Nelson - 2008, 2009
 - Current position: Professor, UC Berkeley
- Arnab Bhattacharyya - 2010
 - Current position: Assistant Professor, National University of Singapore

- Eric Price - 2011
 - Current position: Assistant Professor, UT Austin
- Marco Molinaro - 2012
 - Current position: Assistant Professor, PUC-Rio
- Grigory Yaroslavtsev - 2012
 - Current position: Assistant Professor, Indiana Bloomington
- Yi Li - 2011, 2013
 - Current position: Assistant Professor, Nanyang Technological University
- Huy L. Nguyen - 2013
 - Current position: Assistant Professor, Northeastern
- Michael Crouch – 2014
 - Current position: Bell Labs, Ireland
- Ilya Razenshteyn – 2015
 - Current position: Theory Group at MSR Redmond (currently postdoc at Columbia)
- Zhao Song – 2015, 2016, 2017 (also: winter 2015, winter 2016)
 - Current position: Postdoc at Princeton
- Cameron Musco – 2016
 - Current position: Assistant Professor at UMass Amherst
- Xingguo Li – 2016
 - Current position: Postdoc at Princeton
- Lin Yang – 2017
 - Current position: Assistant Professor at UCLA
- Peilin Zhong – 2017
 - Current position: Graduate Student at Columbia
- Hongyang Zhang – 2017
 - Current position: Postdoc at TTIC
- Chen Shao – 2017
 - Current position: Google

CMU PhD Students

- Ainesh Bakshi (third year)
- Zhili Feng (first year)
- Rajesh Jayaram (third year)
- Praneeth Kacham (first year)
- Hongyang Zhang – graduated – now a postdoc at TTIC (co-advised with Nina Balcan)

CMU Postdocs

- Samson Zhou 2019-21

Visiting Students at CMU

- Amir Zandieh from EPFL: Fall, 2019

Master's/PhD Thesis Committees at CMU

- Hui Han Chan 2017
- Dimitris Konomis 2017-18
- Anit Kumar Sahu 2018
- Colin White 2018
- Vijay Bhattiprolu, 2019
- Hang Liao 2020

Ph.D Thesis Committees Outside of CMU

- Cameron Musco (MIT), 2018
- Ashkay Kamath (UT Austin), 2019

Master's Students Supervised at CMU

- Shuli Jiang 2020
- Dongyu Li 2020

Visiting Undergraduates at CMU

- Peng Ye (Tsinghua), Spring, 2020
- Hanlin Zhu (Tsinghua), Spring, 2020

Undergraduates Advised at CMU

- Manuel Fernandez 2017-19
- Taisuke Yasude 2017-19
- Amulya Musipatla 2018-19
- Arvind Mahankali 2020
- Justin Jia 2020
- Anubhav Baweja 2020

Courses Taught at CMU

- Spring, 2020, Carnegie Mellon University, Algorithms, 15-451, with Danny Sleator
- Fall, 2019, Carnegie Mellon University, Algorithms for Big Data, graduate class 15-859
- Spring, 2019, Carnegie Mellon University, Algorithms, 15-451, with Anupam Gupta
- Spring, 2018, Carnegie Mellon University, Algorithms, 15-451, with Anupam Gupta
- Fall, 2017, Carnegie Mellon University, Algorithms for Big Data, graduate class 15-859

Courses Taught Outside of CMU

- Spring 2 additional courses, 2020, Cathaypath Institute of Science, 8-weeks 3-hr /week, online teaching of an algorithms course
- Spring, 2020, Cathaypath Institute of Science, 8-weeks 3-hr / week, online teaching of an algorithms / algorithms for big data course
- Winter, 2020, Cathaypath Institute of Science, 8-weeks 3-hr / week, online teaching of an algorithms / algorithms for big data course
- Summer, 2019 Cathaypath Institute of Science, 4-week 6-hr / day (4 hours lecturing, 2 hours office hours) teaching of an algorithms / algorithms for big data course in Shanghai
- January 2018, Indian Institute of Technology Kanpur, part of the Global Initiative of Academic Networks (GIAN) program, “Sketching and Sampling for Big Data Analysis”, 14 lectures, each 1 hour
- February 2017, Indian Institute of Technology Kanpur, part of the Global Initiative of Academic Networks (GIAN) program, “Numerical Linear Algebra Methods for Big Data Processing”, 14 lectures, each 1 hour
- Summer, 2016, Shanghai University of Finance and Economics, “Sketching as a Tool for Numerical Linear Algebra”, 16 lectures, each 50-minutes
- Summer, 2015, BASICS Summer School on Communication Complexity, 3 lectures, each 90 minutes
- Summer, 2015, MADALGO Summer School on Streaming Algorithms, 3 lectures, each 90 minutes
- Class project supervisor for Tsinghua Undergraduate Course on Big Data taught by Professor Periklis Papakonstantinou (supervised 6 research projects)

Patents

U.S. Patents issued:

- Server-implemented system and method for providing private inference control
 Patent 8229939 in U.S.
 Co-inventor Jessica Staddon

- Exclusive set system constructions including, but not limited to, applications to broadcast encryption and certificate revocation
 Patent 7818570 in U.S.
 Co-inventors Craig Gentry and Zulfikar Ramzan

- Generation of set coverings with free riders, and generation of ordered sets of meeting points, in systems which include, but are not limited to, systems for broadcast encryption and systems for certificate revocation

Patent 7523304 in U.S.

Co-inventors Craig Gentry and Zulfikar Ramzan

- Random sampling from distributed streams

Patent 8392434 in U.S.

Co-inventor Srikanta Tirthapura

- Aggregate contribution of iceberg queries

Patent 8499003 in U.S.

Co-inventor Jelani Nelson

- Summarizing internet traffic patterns

Patent 8310922 in U.S.

Co-inventor Jelani Nelson

- Computer information retrieval using latent semantic structure via sketches

Patent 8255401 in U.S.

Co-inventor Ken Clarkson

U.S. Patents Filed for:

- By Indiana University Bloomington: Sketching Algorithms for Genomic Data Analysis and Querying in a Secure Enclave, Co-inventors: Can Kockan, Kaiyuan Zhu, Natnatee Dokmai, Nikolai Karpov, M. Oguzhan Kulekci, and S. Cenk Sahinalp

- ARC820150019 Lower-Dimensional Subspace Approximation of a Dataset (previous title: A Method for Subspace Approximation for Minimizing the Maximum Distance)

Co-inventor Kenneth Clarkson

- ARC820130198 merged with ARC820130199: A Fast Method for Regression Using M-Estimators and A Method for Polynomial Kernel Support Vector Machines and Principal Component Analysis

Co-inventors Haim Avron, Ken Clarkson, Huy L. Nguyen

- ARC820130197 A Method and Apparatus for Optimally Finding a CUR Decomposition

Co-inventor Christos Boutsidis

- ARC820120173 Communication and Message-Efficient Protocol for Computing the Intersection Between Two Sets

Co-inventor Grigory Yaroslavtsev

- ARC820120174 Method for Adaptively Breaking Compressed Sensing Schemes and Data

Stream Algorithms

Co-inventor Moritz Hardt

- ARC820130017 Faster Robust Regression Using Exponential Random Variables

Co-inventor Qin Zhang

- ARC820130047 A Method for Managing Deduplication Using Estimated Benefits

Co-inventors David Chambliss, Maohua Lu, M. Corenliu Constantinescu, Joseph S. Glider, Danny Harnik

- ARC820120029 A Method For Fast Distributed Database Frequency Summarization

- ARC820120053 A Method For Estimating the Total Sales over Streaming Bids

Co-inventor Benny Kimelfeld

- ARC820110015 A Method for Efficiently Computing Correlated Aggregates Over a Data Stream

Co-inventor Srikanta Tirthapura

- ARC820110006 A Method for Approximating Klee's Measure Problem and Other Moments on a Stream of Rectangles

Co-inventor Srikanta Tirthapura

- ARC820100022 A Method for Two Parties to Privately Estimate Similarity of Their Datasets

- ARC820100017 Sublinear Time Algorithms for Classification

Co-inventor Ken Clarkson and Elad Hazan

- ARC820080170 Algorithms for Computing Cascaded Aggregates in a Stream

Co-inventor T.S. Jayram

- ARC820070201 An Algorithm for Finding Sparse Directed Spanners, with Applications to Access Control, Networks, and Data Structures

- YOR820130853 Sketching Structured Matrices for Faster Nonlinear Regression

Co-inventors Haim Avron and Vikas Sindhwani

Grants

- National Institutes of Health (NIH) CMU position: \$319,952 (subaward/subcontractor with Indiana Bloomington, PI S. Cenk Sahinalp)
- Office of Naval Research (ONR) CMU position: \$603,654 (PI, co-PI is Jelani Nelson at UC Berkeley)

- National Science Foundation (NSF) Small: \$433,978 (PI)
- Google Faculty Award \$82,535
- Team Member of DARPA Grant Skylark Project on “Randomized Numerical Linear Algebra for Large Scale Data Analysis”. Fully funded 1 day / week for 4 years 2013-2017 (PI: Kenneth Clarkson)
- Received 20K as a subcontractor for Johns Hopkins University for work on an NSF Big Data 2017 grant (PI: Vladimir Braverman)
- Long-term visitor as part of an MOE AcRF Singapore grant at Nanyang Technological Institute (PI: Yi Li)
- Chinese Academy of Sciences PIFI Award – awarded 60,000RMB to visit and collaborate with Professor Xiaoming Sun

Consulting

- Cathaypath Institute of Science 2019-present
 - teaching of algorithms / algorithms for big data courses
 - various hour-long speeches promoting computer science / algorithms in China
- Johns Hopkins University, sub-contracted under Vladimir Braverman – work on streaming algorithms, Fall ‘17-Fall ‘18
- Eureka High School Opportunity Program – mentored 5 high school students interested in machine learning, one-on-one, Fall ‘17 – Fall ‘18

Publications

Electronic copies are available at my DBLP entry:

http://www.informatik.uni-trier.de/~%20ley/pers/hd/w/Woodruff:David_P=.html

Citation information is available on my Google Scholar profile:

<https://scholar.google.com/citations?user=kMmxbblIAAAJ&hl=en>

Books:

David P. Woodruff: Sketching as a Tool for Numerical Linear Algebra, in NOW Publishers Foundations and Trends of Theoretical Computer Science , Vol 10, Issue 1—2, 2014, pp 1—157.

Refereed Conference Publications (Reverse Chronological Order):

172. Debmalya Mandal, Nisarg Shah, and David P. Woodruff:
Optimal Communication-Distortion Tradeoff in Voting, EC, 2020

171. Omri Ben-Eliezer, Rajesh Jayaram, David P. Woodruff, and Eylon Yogev: A Framework for Adversarially Robust Streaming Algorithms, PODS, 2020
170. Sepideh Mahabadi, Ilya Razenshteyn, David P. Woodruff, and Samson Zhou: Non-Adaptive Adaptive Sampling on Turnstile Streams, STOC, 2020
169. Cyrus Rashtchian, Aneesh Sharma, and David P. Woodruff: LSF-Join: Locality Sensitive Filtering for Distributed All-Pairs Similarity Under Skew, WWW (The Web Conference): 2020
168. Praneeth Kacham and David P. Woodruff: Optimal Deterministic Coresets for Ridge Regression, AISTATS, 2020
167. Hang Liao, Barak A. Pearlmutter, Vamsi K. Potluru, and David P. Woodruff: Automatic Differentiation of Sketched Regression, AISTATS, 2020
166. Rajesh Jayram, David P. Woodruff, and Richard Zhang: Span Recovery for Deep Neural Networks with Applications to Input Obfuscation, ICLR, 2020
165. Tanqiu Jiang, Yi Li, Honghao Lin, Yisong Ruan, and David P. Woodruff: Learning-Augmented Data Stream Algorithms, ICLR, 2020
164. Manuel Fernandez, David P. Woodruff, and Taisuke Yasuda: Graph Spanners in the Message-Passing Model, ITCS, 2020
163. Shafi Goldwasser, Ofer Grossman, Sidhanth Mohanty, and David P. Woodruff: Pseudo-deterministic Streaming, ITCS, 2020
162. Thomas D. Ahle, Michael Kapralov, Jakob B. T. Knudsen, Rasmus Pagh, Ameya Velingker, David P. Woodruff, and Amir Zandieh: Oblivious Sketching of High-Degree Polynomial Kernels, SODA, 2020
161. Yi Li, Ruosong Wang, and David P. Woodruff: Tight Bounds for the Subspace Sketch Problem, SODA, 2020
160. Santosh Vempala, Ruosong Wang, and David P. Woodruff: The Communication Complexity of Optimization, SODA, 2020.
- Invited to Highlights of Algorithms, 2020.
159. Debmalya Mandal, Ariel Procaccia, Nisarg Shah, and David P. Woodruff: Efficient and Thrifty Voting by Any Means Necessary, NeurIPS, 2019.

Selected for Oral Presentation (1 of 36 out of 1428 papers)

158. Frank Ban, David P. Woodruff, and Richard Zhang: Regeularized Weighted Low Rank Approximation, NeurIPS, 2019
157. Huaian Diao, Rajesh Jayaram, Zhao Song, Wen Sun, and David P. Woodruff: Optimal Sketching for Kronecker Product Regression and Low Rank Approximation, NeurIPS, 2019

156. Huaian Diao, Zhao Song, David P. Woodruff, and Xin Yang: Total Least Squares in Input Sparsity Time, NeurIPS, 2019
155. Michela Meister, Tamas Sarlos, and David P. Woodruff: Tight Dimensionality Reduction for Sketching Low Degree Polynomial Kernels, NeurIPS, 2019
154. Zhao Song, David P. Woodruff, and Peilin Zhong: Average Case Column Subset Selection for Entrywise ℓ_1 -Norm, NeurIPS, 2019
153. Zhao Song, David P. Woodruff, and Peilin Zhong: Towards a Zero-One Law for Column Subset Selection, NeurIPS, 2019
- 15.2 Manuel Fernandez, David P. Woodruff, and Taisuke Yasude: The Query Complexity of Mastermind with L_p Distances, APPROX, 2019
151. Rajesh Jayaram and David P. Woodruff: Towards Optimal Moment Estimation in Streaming and Distributed Models, APPROX, 2019
150. Kenneth L. Clarkson, Ruosong Wang, and David P. Woodruff: Dimensionality Reduction for Tukey Regression, ICML, 2019
149. Manuel Fernandez, David P. Woodruff, and Taisuke Yasude: Tight Kernel Query Complexity of Kernel Ridge Regression and Kernel k -Means Clustering, ICML, 2019. Selected for a long talk.
148. Ravi Kumar, Rina Panigrahy, Ali Rahimi, and David P. Woodruff: Faster Algorithms for Boolean Matrix Factorization, ICML, 2019. Selected for a long talk.
147. Pranjal Awasthi, Ainesh Bakshi, Nina Balcan, Colin White, and David P. Woodruff: Robust Communication-Optimal Distributed Clustering, ICALP, 2019
146. Xiaoming Sun, David P. Woodruff, Guang Yang, and Jialin Zhang: Querying a Matrix through Matrix-Vector Products, ICALP, 2019
145. David P. Woodruff and Guang Yang: Separating k -Player from t -Player One-Way Communication, ICALP, 2019
144. Ainesh Bakshi, Rajesh Jayaram, and David P. Woodruff: Learning Two Layer Rectified Neural Networks in Polynomial Time, COLT, 2019
143. Yu Cheng, Ilias Diakonikolas, Rong Ge, and David P. Woodruff: Faster Algorithms for High Dimensional Robust Covariance Estimation, COLT, 2019
142. Piotr Indyk, Ali Vakilian, Tal Wagner, and David P. Woodruff: Sample-Optimal Low-Rank Approximation of Distance Matrices, COLT, 2019
141. Rajesh Jayram, Gokarna Sharma, Srikanta Tirthapura, and David P. Woodruff: Weighted Reservoir Sampling from Distributed Streams, PODS, 2019
140. Vladimir Braverman, Moses Charikar, William Kuszmaul, David P. Woodruff, and Lin F. Yang: The One-Way communication Complexity of Dynamic Time Warping, SOCG, 2019,

invited to the special issue for SOCG, 2019

139. John Hainline, Brendan Juba, Hai S. Le, and David P. Woodruff: Conditional Sparse L_p -norm Regression with Optimal Probability, AISTATS, 2019

138. Can Kockan, Kaiyuan Zhu, Natnatee Dokmai, Nikolai Karpov, M. Oguzhan Kulecki, David P. Woodruff, and S. Cenk Sahinalp: Sketching Algorithms for Genomic Data Analysis and Querying in a Secure Enclave, RECOMB, 2019. Full version submitted to Nature Methods.

137. Xiaofei Shi and David P. Woodruff: Sublinear Time Numerical Linear Algebra for Structured Matrices, AAI, 2019

136. Nina Balcan, Yi Li, David P. Woodruff, and Hongyang Zhang: Testing Matrix Rank, optimally, SODA, 2019

135. Frank Ban, Vijay Bhattiprolu, Karl Bringmann, Pavel Kolev, Euiwoong Lee, and David P. Woodruff: A PTAS for ℓ_p Low Rank Approximation, SODA, 2019

134. Zhao Song, David P. Woodruff, and Peilin Zhong: Relative Error Tensor Low Rank Approximation, SODA, 2019

133. Ruosong Wang and David P. Woodruff: Tight Bounds for ℓ_p Oblivious Subspace Embeddings, SODA, 2019, invited to the special issue for SODA, 2019

132. Ainesh Bakshi and David P. Woodruff: Sublinear Time Low-Rank Approximation of Distance Matrices, NeurIPS, 2018, selected for spotlight presentation

131. Alexander Munteanu, Chris Schweigelshohn, Christian Sohler, and David P. Woodruff: On Coresets for Logistic Regression, NeurIPS, 2018, selected for spotlight presentation. Also presented at Informatik, 2019

130. Roie Levin, Anish Sevekari, and David P. Woodruff: Robust Subspace Approximation in a Stream, NeurIPS, 2018, selected for spotlight presentation

129. Christian Sohler and David P. Woodruff: Strong Coresets for k-Median and Subspace Approximation, Goodbye Dimension, FOCS, 2018

128. Rajesh Jayaram and David P. Woodruff: Perfect L_p Sampling in a Data Stream, FOCS, 2018.

127. Aditya Krishnan, Sidhant Mohanty, and David P. Woodruff: On Sketching the q to p Norms, APPROX, 2018.

126. Yi Li, Vasileios Nakos, and David P. Woodruff: On Low-Risk Heavy Hitters and Sparse Recovery Schemes, APPROX, 2018

125. Graham Cormode, Charlie Dickens, and David P. Woodruff: Leveraging Well-Conditioned Bases: Streaming and Distributed Summaries in Minkowski p -Norms, ICML, 2018. Selected for a long talk.

124. Vladimir Braverman, Stephen R. Chestnut, Robert Krauthgamer, Yi Li, David P. Woodruff,

and Lin Yang: Matrix Norms in Data Streams: Faster, Multi-Pass, and Row-Order, ICML, 2018

123. Yogesh Dahiya, Dimitris Konomis, and David P. Woodruff: An Empirical Evaluation of Sketching for Numerical Linear Algebra, KDD, 2018

122. Vladimir Braverman, Emanuele Viola, David P. Woodruff, and Lin Yang: Revisiting Frequency Moment Estimation in Random Order Streams, ICALP, 2018

121. Sumit Ganguly and David P. Woodruff: High Probability Frequency Moment Sketches, ICALP, 2018

120. Vasileios Nakos, Xiaofei Shi, David P. Woodruff, and Hongyang Zhang: Improved Algorithms for Adaptive Compressed Sensing, ICALP, 2018

119. Rajesh Jayaram and David P. Woodruff: Data Streams with Bounded Deletions, PODS, 2018

118. Huaian Diao, Zhao Song, Wen Sun, and David P. Woodruff: Sketching for Kronecker Product Regression and P-Splines, AISTATS, 2018. Selected for oral presentation

117. David P. Woodruff and Qin Zhang: Distributed Statistical Estimation of Matrix Products, PODS, 2018

116. Cameron Musco, Praneeth Netrapalli, Aaron Sidford, Shashanka Ubaru, and David P. Woodruff: Spectrum Approximation Beyond Fast Matrix Multiplication, ITCS, 2018

115. Nina Balcan, Yingyu Liang, David P. Woodruff, and Hongyang Zhang: Matrix Completion and Related Problems via Strong Duality, ITCS, 2018

114. Cameron Musco and David P. Woodruff: Is Input Sparsity Time Possible for Kernel Low-Rank Approximation?, NIPS, 2017

113. Jarvis Haupt, Xingguo Li, and David P. Woodruff: Near Optimal Sketching of Low-Rank Tensor Regression, NIPS, 2017

112. Karl Bringmann, Pavel Kolev, and David P. Woodruff: Approximation Algorithms for ℓ_0 -Low Rank Approximation, NIPS 2017

111. Michael Kapralov, Jelani Nelson, Jakub Pachocki, Zhengyu Wang, David P. Woodruff, and Mobin Yahyazadeh: Optimal Lower Bounds for Universal Relation, and for Samplers and Finding Duplicates in Streams, FOCS, 2017

110. Cameron Musco and David P. Woodruff: Sublinear Time Low-Rank Approximation of Positive Semidefinite Matrices, FOCS, 2017

110. Haim Avron and Kenneth L. Clarkson: Sharper Bounds for Regularized Data Fitting, RANDOM, 2017

109. Flavio Chierichetti, Sreenivas Gollapudi, Ravi Kumar, Silvio Latanzi, Rina Panigrahy, and David P. Woodruff: Algorithms for ℓ_p Low-Rank Approximation, ICML, 2017

108. Yi Li and David P. Woodruff: Embeddings of Schatten Norms with Applications to Data Streams, ICALP, 2017
107. Eric Price, Zhao Song, and David P. Woodruff: Fast Regression with an ℓ_∞ Guarantee, ICALP, 2017
106. Zhao Song, David P. Woodruff, and Peilin Zhong: Low Rank Approximation with Entrywise ℓ_1 -Norm Error, STOC, 2017
105. Vladimir Braverman, Stephen R. Chestnut, Nikita Ivkin, Jelani Nelson, Zhengyu Wang, and David P. Woodruff: BPTree: an L2 Heavy Hitters Algorithm Using Constant Memory, PODS, 2017
104. Kenneth L. Clarkson and David P. Woodruff: Low-Rank PSD Approximation in Input-Sparsity Time, SODA, 2017
103. Santosh Vempala and David P. Woodruff: Adaptive Matrix Vector Product, SODA, 2017
102. Jiecao Chen, He Sun, David P. Woodruff, and Qin Zhang: Communication-Optimal Distributed Clustering, NIPS, 2016
101. Zhao Song, David P. Woodruff, and Huan Zhang: Sublinear Time Orthogonal Tensor Decomposition, NIPS, 2016
100. Hossein Esfandiari, Mohammad Taghi Hajiaghayi, and David P. Woodruff: Applications of Uniform Sampling: Densest Subgraph and Beyond, SPAA, 2016
99. Yi Li and David P. Woodruff: Tight Bounds for Sketching the Operator Norm, Schatten Norms, and Subspace Embeddings, RANDOM, 2016
98. Michael Crouch, Andrew McGregor, Gregory Valiant, and David P. Woodruff: Stochastic Streams: Sample Complexity vs. Space Complexity, ESA, 2016
97. Maria-Florina Balcan, Yingyu Liang, Le Song, David P. Woodruff, and Bo Xie: Communication Efficient Distributed Kernel Principal Component Analysis, KDD, 2016
96. Michael Kapralov, Vamsi k. Potluru, and David P. Woodruff: How to Fake Multiply by a Gaussian Matrix, ICML, 2016
95. Michael B. Cohen, Jelani Nelson, and David P. Woodruff: Optimal Approximate Matrix Product in Terms of Stable Rank, ICALP, 2016
94. David P. Woodruff: New Algorithms for Heavy Hitters in Data Streams, ICDT (invited paper), 2016
93. Arnab Bhattacharyya, Palash Dey, and David P. Woodruff: An Optimal Algorithm for l_1 -Heavy Hitters in Insertion Streams and Related Problems, PODS, 2016
92. Vladimir Braverman, Stephen R. Chestnut, Lin F. Yang, and David P. Woodruff: Streaming Space Complexity of Nearly All Functions of One Variable on Frequency Vectors, PODS, 2016

91. Yuqing Ai, Wei Hu, Yi Li, and David P. Woodruff: New Characterizations in Turnstile Streams with Applications, CCC, 2016
90. Vladimir Braverman, Stephen R. Chestnut, Nikita Ivkin, and David P. Woodruff: Beating CountSketch for Heavy Hitters in Insertion Streams, STOC, 2016
89. Yi Li and David P. Woodruff: On Approximating Functions of the Singular Values in a Stream, STOC, 2016
88. Mark Braverman, Ankit Garg, Tengyu Ma, Huy L. Nguyen, and David P. Woodruff: Communication Lower Bounds for Statistical Estimation Problems via a Distributed Data Processing Inequality, STOC, 2016
87. Christos Boutsidis, David P. Woodruff, and Peilin Zhong: Optimal Principal Component Analysis in Distributed and Streaming Models, STOC, 2016
86. Ilya Razenshteyn, Zhao Song, and David P. Woodruff: Weighted Low Rank Approximations with Provable Guarantees, STOC, 2016
85. David P. Woodruff, Peilin Zhong: Distributed Low Rank Approximation of Implicit Functions of a Matrix, ICDE, 2016
84. Arturs Backurs, Piotr Indyk, Ilya Razenshteyn, and David P. Woodruff: Nearly-optimal Bounds for Sparse Recovery in Generic Norms, with Applications to k-median Sketching, SODA, 2016
83. Alexandr Andoni, Jiecao Chen, Bo Qin, Robert Krauthgamer, David P. Woodruff, and Qin Zhang: On Sketching Quadratic Forms, ITCS, 2016
82. Xiaoming Sun, David P. Woodruff: Tight Bounds for Graphs Problems in Insertion Streams, APPROX, 2015
81. Ken Clarkson, David P. Woodruff: Input Sparsity and Hardness for Robust Subspace Approximation. FOCS, 2015
80. Dirk van Gucht, Ryan Williams, David P. Woodruff, Qin Zhang: The Communication Complexity of Distributed Set-Joins with Applications to Matrix Multiplication, PODS, 2015
79. Marco Molinaro, David P. Woodruff, Grigory Yaroslavtsev: Amplification of One-Way Information Complexity via Codes and Noise Sensitivity, ICALP, 2015
78. Omri Weinstein, David P. Woodruff: The Simultaneous Communication of Disjointness with Applications to Data Streams, ICALP 2015
77. Ken Clarkson, David P. Woodruff: Sketching for M-Estimators: A Unified Approach to Robust Regression. SODA, 2015
76. Haim Avron, Huy L. Nguyen, David P. Woodruff: Subspace Embeddings for the Polynomial Kernel. NIPS, 2014

75. David P. Woodruff: Low Rank Approximation Lower Bounds in Row-Update Streams. NIPS, 2014
74. Nina Balcan, Vandana Kanchanapally, Yingyu Liang, David P. Woodruff: Improved Distributed Principal Component Analysis. NIPS, 2014
73. Ravi Kannan, Santosh Vempala, David P. Woodruff: Principal Component Analysis and Higher Correlations for Distributed Data. COLT, 2014: 1040-1057
72. Joshua Brody, Amit Chakrabarti, Ranganath Kondapally, David P. Woodruff, Grigory Yaroslavtsev: Certifying Equality with Limited Interaction. APPROX-RANDOM, 2014: 545-581
71. Yi Li, Zhengyu Wang, David P. Woodruff: Improved Testing of Low Rank Matrices. KDD, 2014: 691-700, invited to the special issue in ACM Transactions on Knowledge Discovery from Data for select papers from KDD, 2014
70. Joshua Brody, Amit Chakrabarti, Ranganath Kondapally, David P. Woodruff, Grigory Yaroslavtsev: Beyond Set Disjointness: the Communication Complexity of Finding the Intersection. PODC, 2014: 106-113
69. Michael Kapralov, David P. Woodruff: Spanners and Sparsifiers in Dynamic Streams. PODC, 2014: 272-281, invited to the special issue in Distributed Computing for select papers from PODC, 2014
68. Rasmus Pagh, Morten Stockel, David P. Woodruff: Is Min-Wise Hashing Optimal for Summarizing Set Intersection? PODS, 2014: 109-120
67. Yi Li, Xiaoming Sun, Chengu Wang, David P. Woodruff: On the Communication Complexity of Linear Algebraic Problems in the Message Passing Model. DISC, 2014
66. Christos Boutsidis, David P. Woodruff: Optimal CUR Matrix Decompositions. STOC, 2014: 353-362
65. Yi Li, Huy L. Nguyen, David P. Woodruff: Turnstile Streaming Algorithms Might as Well be Linear Sketches. STOC, 2014: 174-183
64. Benny Kimelfeld, Jan Vondrak, David P. Woodruff: Multi-Tuple Deletion Propagation: Approximations and Complexity. VLDB 2014: 1558-1569
63. David P. Woodruff, Qin Zhang: An Optimal Lower Bound for Distinct Elements in the Message Passing Model. SODA 2014: 718-733
62. Yi Li, Huy L. Nguyen, David P. Woodruff: On Sketching Matrix Norms and the Top Singular Vector. SODA 2014: 1562-1581

61. Yi Li, David P. Woodruff: A Tight Lower Bound for High Frequency Moment Estimation with Small Error. APPROX-RANDOM 2013: 623-638
60. David P. Woodruff, Qin Zhang: Subspace Embeddings and Lp-Regression Using Exponential Random Variables. COLT 2013: 546-567
59. Kenneth L. Clarkson, Petros Drineas, Malik Magdon-Ismail, Michael W. Mahoney, Xiangrui Meng, David P. Woodruff: The Fast Cauchy Transform and Faster Robust Linear Regression. SODA 2013: 466-477
58. Eric Price, David P. Woodruff: Lower Bounds for Adaptive Sparse Recovery. SODA 2013: 652-663
57. Marco Molinaro, David P. Woodruff, Grigory Yaroslavtsev: Beating the Direct Sum Theorem in Communication Complexity with Implications for Sketching. SODA 2013: 1738-1756
56. Kenneth L. Clarkson, David P. Woodruff: Low Rank Approximation and Regression in Input Sparsity Time. STOC 2013: 81-90, invited to the Journal of the ACM
55. Moritz Hardt, David P. Woodruff: How Robust are Linear Sketches to Adaptive Inputs? STOC 2013: 121-130
54. David P. Woodruff, Qin Zhang: When Distributed Computation Is Communication Expensive. DISC 2013: 16-30, invited to the special issue in Distributed Computing for select papers from DISC, 2013
53. Haim Avron, Vikas Sindhwani, David P. Woodruff: Sketching Structured Matrices for Faster Nonlinear Regression. NIPS 2013
52. Anna C. Gilbert, Brett Hemenway, Martin J. Strauss, David P. Woodruff, Mary Wootters: Reusable Low-error Compressive Sampling Schemes Through Privacy. SSP 2012.
51. Michael W. Mahoney, Petros Drineas, Malik Magdon-Ismail, David P. Woodruff: Fast Approximation of Matrix Coherence and Statistical Leverage. ICML 2012
50. Jelani Nelson, Huy L. Nguyn, David P. Woodruff: On Deterministic Sketching and Streaming for Sparse Recovery and Norm Estimation. APPROX-RANDOM 2012: 627-638
49. Srikanta Tirthapura, David P. Woodruff: A General Method for Estimating Correlated Aggregates over a Data Stream. ICDE 2012: 162-173
48. Eric Price, David P. Woodruff: Applications of the Shannon-Hartley Theorem to Data Streams and Sparse Recovery. ISIT 2012: 2446-2450

47. Andrew McGregor, A. Pavan, Srikanta Tirthapura, David P. Woodruff: Space-efficient Estimation of Statistics over Sub-sampled Streams. PODS 2012: 273-282
46. Srikanta Tirthapura, David P. Woodruff: Rectangle-efficient Aggregation in Spatial Data Streams. PODS 2012: 283-294
45. David P. Woodruff, Qin Zhang: Tight Bounds for Distributed Functional Monitoring. STOC 2012: 941-960
44. Joshua Brody, David P. Woodruff: Streaming Algorithms with One-Sided Estimation. APPROX-RANDOM 2011: 436-447
43. Rolf Klein, Rainer Penninger, Christian Sohler, David P. Woodruff: Tolerant Algorithms. ESA 2011: 736-747
42. Piotr Indyk, Eric Price, David P. Woodruff: On the Power of Adaptivity in Sparse Recovery. FOCS 2011: 285-294
41. Eric Price, David P. Woodruff: $(1 + \epsilon)$ -Approximate Sparse Recovery. FOCS 2011: 295-304, invited to the special issue in Algorithmica on group testing and compressed sensing
40. Piotr Berman, Arnab Bhattacharyya, Elena Grigorescu, Sofya Raskhodnikova, David P. Woodruff, Grigory Yaroslavtsev: Steiner Transitive-Closure Spanners of Low-Dimensional Posets. ICALP (1) 2011: 760-772
39. Arnab Bhattacharyya, Piotr Indyk, David P. Woodruff, Ning Xie: The Complexity of Linear Dependence Problems in Vector Spaces. ICS 2011: 496-508
38. T. S. Jayram, David P. Woodruff: Optimal Bounds for Johnson-Lindenstrauss Transforms and Streaming Problems with Sub-Constant Error. SODA 2011: 1-10, invited to the special issue in Transactions on Algorithms for select papers from SODA, 2011
37. David P. Woodruff: Near-optimal Private Approximation Protocols via a Black Box transformation. STOC 2011: 735-744
36. Daniel M. Kane, Jelani Nelson, Ely Porat, David P. Woodruff: Fast Moment Estimation in Data Streams in Optimal Space. STOC 2011: 745-754
35. Christian Sohler, David P. Woodruff: Subspace Embeddings for the L1-norm with Applications. STOC 2011: 755-764
34. Srikanta Tirthapura, David P. Woodruff: Optimal Random Sampling from Distributed Streams Revisited. DISC 2011: 283-297
33. Arnab Bhattacharyya, Elena Grigorescu, Madhav Jha, Kyomin Jung, Sofya Raskhodnikova, David P. Woodruff: Lower Bounds for Local Monotonicity Reconstruction from

Transitive-Closure Spanners. APPROX-RANDOM 2010: 448-461

32. David P. Woodruff: A Quadratic Lower Bound for Three-Query Linear Locally Decodable Codes over Any Field. APPROX-RANDOM 2010: 766-779

31. Kenneth L. Clarkson, Elad Hazan, David P. Woodruff: Sublinear Optimization for Machine Learning. FOCS 2010: 449-457

30. David P. Woodruff: Additive Spanners in Nearly Quadratic Time. ICALP (1) 2010: 463-474

29. Daniel M. Kane, Jelani Nelson, David P. Woodruff: An Optimal Algorithm for the Distinct Elements Problem. PODS 2010: 41-52, invited to the Journal of the ACM

28. Jelani Nelson, David P. Woodruff: Fast Manhattan Sketches in Data Streams. PODS 2010: 99-110

27. Dan Feldman, Morteza Monemizadeh, Christian Sohler, David P. Woodruff: Coresets and Sketches for High Dimensional Subspace Approximation Problems. SODA 2010: 630-649

26. Morteza Monemizadeh, David P. Woodruff: 1-Pass Relative-Error Lp-Sampling with Applications. SODA 2010: 1143-1160

25. Daniel M. Kane, Jelani Nelson, David P. Woodruff: On the Exact Space Complexity of Sketching and Streaming Small Norms. SODA 2010: 1161-1178

24. Khanh Do Ba, Piotr Indyk, Eric Price, David P. Woodruff: Lower Bounds for Sparse Recovery. SODA 2010: 1190-1197

23. Alexandr Andoni, Khanh Do Ba, Piotr Indyk, David P. Woodruff: Efficient Sketches for Earth-Mover Distance, with Applications. FOCS 2009: 324-330

22. T. S. Jayram, David P. Woodruff: The Data Stream Space Complexity of Cascaded Norms. FOCS 2009: 765-774

21. David P. Woodruff: The Average-case Complexity of Counting Distinct Elements. ICDT 2009: 284-295

20. Arnab Bhattacharyya, Elena Grigorescu, Kyomin Jung, Sofya Raskhodnikova, David P. Woodruff: Transitive-closure Spanners. SODA 2009: 932-941

19. Kenneth L. Clarkson, David P. Woodruff: Numerical Linear Algebra in the Streaming Model. STOC 2009: 205-214

18. David P. Woodruff: Corruption and Recovery-Efficient Locally Decodable Codes. APPROX-RANDOM 2008: 584-595

17. Alexandre V. Evfimievski, Ronald Fagin, David P. Woodruff: Epistemic Privacy. PODS 2008: 171-180
16. David P. Woodruff: Revisiting the Efficiency of Malicious Two-Party Computation. EUROCRYPT 2007: 79-96
15. Xiaoming Sun, David P. Woodruff: The Communication and Streaming Complexity of Computing the Longest Common and Increasing Subsequences. SODA 2007: 336-345
14. David P. Woodruff: Better Approximations for the Minimum Common Integer Partition Problem. APPROX-RANDOM 2006: 248-259
13. Zulfikar Ramzan, David P. Woodruff: Fast Algorithms for the Free Riders Problem in Broadcast Encryption. CRYPTO 2006: 308-325
12. Craig Gentry, Zulfikar Ramzan, David P. Woodruff: Explicit Exclusive Set Systems with Applications to Broadcast Encryption. FOCS 2006: 27-38
11. David P. Woodruff: Lower Bounds for Additive Spanners, Emulators, and More. FOCS 2006: 389-398
10. Piotr Indyk, David P. Woodruff: Polylogarithmic Private Approximations and Efficient Matching. TCC 2006: 245-264
9. David P. Woodruff, Sergey Yekhanin: A Geometric Approach to Information-Theoretic Private Information Retrieval. IEEE Conference on Computational Complexity 2005: 275-284
8. Marten van Dijk, Robert Granger, Dan Page, Karl Rubin, Alice Silverberg, Martijn Stam, David P. Woodruff: Practical Cryptography in High Dimensional Tori. EUROCRYPT 2005: 234-250
7. Piotr Indyk, David P. Woodruff: Optimal Approximations of the Frequency Moments of Data Streams. STOC 2005: 202-208
6. David P. Woodruff, Jessica Staddon: Private Inference Control. ACM Conference on Computer and Communications Security 2004: 188-197
5. Marten van Dijk, David P. Woodruff: Asymptotically Optimal Communication for Torus-Based Cryptography. CRYPTO 2004: 157-178
4. Hanson Zhou, David P. Woodruff: Clustering via Matrix Powering. PODS 2004: 136-142
3. David P. Woodruff: Optimal Space Lower Bounds for all Frequency Moments. SODA 2004: 167-175
2. Piotr Indyk, David P. Woodruff: Tight Lower Bounds for the Distinct Elements Problem.

FOCS 2003: 283-288

1. David P. Woodruff, Marten van Dijk: Cryptography in an Unbounded Computational Model. EUROCRYPT 2002: 149-164

Refereed Journal Publications:

30. Can Kockan, Kaiyuan Zhu, Natnatee Dokmai, Nikolai Karpov, M. Oguzhan Kulecki, David P. Woodruff, and S. Cenk Sahinalp: Sketching Algorithms for Genomic Data Analysis and Querying in a Secure Enclave, Nature Methods, to appear

29. Yi Li, Huy Le Nguyen, and David P. Woodruff: On Approximating Matrix Norms in a Stream, SICOMP, to appear

28. Nina Balcan, Yingyu Liang, Zhao Song, David P. Woodruff, and Hongyang Zhang: Non-Convex Matrix Completion and Related Problems via Strong Duality, Journal of Machine Learning Research, 2019

27. Arnab Bhattacharyya, Palash Dey, and David P. Woodruff: An Optimal Algorithm for ℓ_1 Heavy Hitters in Insertion Streams and Related Problems, ACM Transactions on Algorithms, 2019

26. Alexandr Andoni, Jiecao Chen, Robert Krauthgamer, Bo Qin, David P. Woodruff, Qin Zhang: On Sketching Quadratic Forms, submitted to Journal of the ACM, 2017 (passed initial screening, revising according to reviewer comments).

25. Kenneth L. Clarkson, David P. Woodruff: Low Rank Approximation and Regression in Input Sparsity Time, Journal of the ACM, 2017.

24. David P. Woodruff, Qin Zhang: When Distributed Computation is Communication Expensive, Distributed Computing, 30 (5), 309-323, 2017

23. Haim Avron, Kenneth L. Clarkson, David P. Woodruff: Faster Kernel Ridge Regression Using Sketching and Preconditioning, accepted to SIMAX with minor revisions

22. Christos Boutsidis, David P. Woodruff: Optimal CUR Matrix Decompositions, SICOMP 46(2): 543-589, 2017

21. [Srikanta Tirthapura](#), David P. Woodruff: A General Method for Estimating Correlated Aggregates Over a Data Stream. [Algorithmica 73\(2\)](#): 235-260, 2015

20. [Yung-Yu Chung](#), [Srikanta Tirthapura](#), David P. Woodruff: A Simple Message-Optimal Algorithm for Random Sampling from a Distributed Stream. [IEEE Trans. Knowl. Data Eng. 28\(6\)](#): 1356-1368, 2016

19. [Yuval Rabani](#), [Andréa W. Richa](#), [Jared Saia](#), David P. Woodruff: Editorial to the Special Issue on SODA'12. [ACM Trans. Algorithms 12\(1\)](#): 1:1, 2016

18. [Mina Ghashami](#), [Edo Liberty](#), [Jeff M. Phillips](#), David P. Woodruff: Frequent Directions: Simple and Deterministic Matrix Sketching. [SIAM J. Comput. 45\(5\)](#): 1762-1792, 2016
17. [Kenneth L. Clarkson](#), [Petros Drineas](#), [Malik Magdon-Ismail](#), [Michael W. Mahoney](#), [Xiangrui Meng](#), David P. Woodruff: The Fast Cauchy Transform and Faster Robust Linear Regression. [SIAM J. Comput. 45\(3\)](#): 763-810, 2016
16. [Mark Braverman](#), David P. Woodruff: Guest Editorial for Information Complexity and Applications. [Algorithmica 76\(3\)](#): 595-596, 2016
15. [Joshua Brody](#), [Amit Chakrabarti](#), [Ranganath Kondapally](#), David P. Woodruff, [Grigory Yaroslavtsev](#): Certifying Equality With Limited Interaction. [Algorithmica 76\(3\)](#): 796-845, 2016
14. Periklis Papakonstantinou, David P. Woodruff, and Guang Yang: True Randomness from Big Data, Scientific Reports, 2016
13. Andrew McGregor, Aduri Pavan, Srikanta Tirthapura, David P. Woodruff: Space-Efficient Estimation of Statistics over Sub-Sampled Streams. *Algorithmica*, 2016
12. David P. Woodruff: Data Streams and Applications in Computer Science 2014: Bulletin of the EATCS 114.
11. Piotr Berman, Arnab Bhattacharyya, Elena Grigorescu, Sofya Raskhodnikova, David Woodruff, Grigory Yaroslavtsev: Steiner Transitive-Closure Spanners of Low-Dimensional Posets. *Combinatorica*, 2014.
10. Jelani Nelson, Huy L. Nguyen, David P. Woodruff: On Deterministic Sketching and Streaming for Sparse Recovery and Norm Estimation. *Linear Algebra and Applications*, March 2013.
9. Benny Kimelfeld, Jan Vondrak, David P. Woodruff: Multi-Tuple Deletion Propagation: Approximations and Complexity. *PVLDB 6(13)*: 1558-1569 (2013)
8. T. S. Jayram, David P. Woodruff: Optimal Bounds for Johnson-Lindenstrauss Transforms and Streaming Problems with Subconstant Error. *ACM Transactions on Algorithms 9(3)*: 26 (2013)
7. Petros Drineas, Malik Magdon-Ismail, Michael W. Mahoney, David P. Woodruff: Fast Approximation of Matrix Coherence and Statistical Leverage. *Journal of Machine Learning (JMLR)*, volume 13, 2012.
6. Kenneth L. Clarkson, Elad Hazan, David P. Woodruff: Sublinear Optimization for Machine Learning. *J. ACM 59(5)*: 23 (2012)
5. Arnab Bhattacharyya, Elena Grigorescu, Kyomin Jung, Sofya Raskhodnikova, David P. Woodruff: Transitive-Closure Spanners. *SIAM J. Comput. 41(6)*: 1380-1425 (2012)

4. Arnab Bhattacharyya, Elena Grigorescu, Madhav Jha, Kyomin Jung, Sofya Raskhodnikova, David P. Woodruff: Lower Bounds for Local Monotonicity Reconstruction from Transitive-Closure Spanners. *SIAM J. Discrete Math.* 26(2): 618-646 (2012)
3. David P. Woodruff: A Quadratic Lower Bound for Three-Query Linear Locally Decodable Codes over Any Field. *J. Comput. Sci. Technol.* 27(4): 678-686 (2012)
2. Alexandre V. Evfimievski, Ronald Fagin, David P. Woodruff: Epistemic Privacy. *J. ACM* 58(1): 2 (2010)
1. David P. Woodruff, Sergey Yekhanin: A Geometric Approach to Information-Theoretic Private Information Retrieval. *SIAM J. Comput.* 37(4): 1046-1056 (2007)

Short Surveys, Theses:

4. Ilias Diakonikolas, Santosh Vempala, David P. Woodruff: Algorithmic High-Dimensional Robust Statistics, Simons Vignette, 2019
3. David P. Woodruff: Frequency Moments. *Encyclopedia of Database Systems 2009*: 1169-1170
2. David P. Woodruff: Efficient and Private Distance Approximation in the Communication and Streaming Models. MIT Ph.D. Thesis, 2007
1. David P. Woodruff: Cryptography in an Unbounded Computational Model. MIT M.Eng. Thesis, 2002

Invited Plenary talks at Conferences

- ICDT 2016
- ESA 2017
- AAAC 2018
- AAIM 2019
- CSR 2019

Invited Talks at Workshops

96. Data Stream Algorithms Workshop at Institute for Mathematical Sciences at Nanyang Technological Institute, Singapore (upcoming)
95. Information-Theoretic Methods in Complexity Theory, Princeton, 2020 (upcoming)

94. Open Data Science Conference, ODSC, Boston, 2020
93. SIAM Conference on Optimization, 2020 (upcoming)
92. Theory Underlying Algorithms, TUNGA, 2020
91. Information Theory and Applications, 2020
90. Israel Theory Day, Open University of Israel, 2019
89. Theory of Data Science and Deep Learning program, Tel-Aviv and University of Haifa, 2019
88. Randomized Numerical Linear Algebra and Applications in Data Science, University of Newcastle, Australia, 2019
87. Simons Foundations of Data Science Reunion Workshop, 2019
86. Session on Semidefinite Programming and Low Rank Approximation, INFORMS, 2019
85. Workshop on Bridging Discrete and Continuous Optimization, DIMACS, 2019
84. Workshop on Recent Trends in Clustering and Classification, TTIC, 2019
83. Workshop at STOC on Data Science Through a Geometric Lens, 2019
82. Metric Embeddings and Dimension Reduction, Northwestern University, 2019
81. Simons Foundations of Data Science Boot Camp, 2018
80. International Symposium on Mathematical Programming, University of Bordeaux, 2018
79. International Symposium on Big Data, Tianjin University of Technology, 2018
78. Sublinear Algorithms Workshop, MIT, 2018
77. Workshop on High-Dimensional Robust Statistics, TTI-Chicago, 2018
76. Workshop on Data Summarization, University of Warwick, 2018
75. Mathematical and Computational Challenges in Real-Time Decision Making at Simons, 2018
74. Workshop on Algorithms and Randomness at Georgia Tech, 2018
73. Workshop on Theoretical Computer Science and Statistical Physics Methods in Machine Learning in Bangalore, 2017

72. Workshop at FOCS on Sketching as a Tool for Everything, 2017
71. Allerton Session on Statistical Inference and Machine Learning, 2017
70. Simons Workshop on Fast Iterative Methods in Optimization, 2017
69. Simons Collaboration on Algorithms and Geometry, 2017
68. SF Data Institute Conference, 2017
67. Shonan Workshop on Processing Big Data Streams, 2017
66. Workshop on Algorithmic Challenges in Data Science at EPFL, 2017
65. Toca Theory Day at Google, "Relative Error Tensor Approximation", 2017
64. Dagstuhl Computational Geometry workshop, long survey on "Sketching as a Tool for Geometric Problems", 2017
63. Bellairs Workshop on Data, Learning, and Optimization, long survey on "Parameterized Complexity of Matrix Factorization Problems", 2017
62. Bellairs Workshop on Data, Learning, and Optimization, long survey on "Sketching as a Tool for Numerical Linear Algebra", 2017
61. Banff Communication Complexity and Applications II, "Sketching and Streaming Matrix Norms", 2017
60. Dagstuhl Parameterized Complexity workshop, long survey on "Parameterized Complexity of Matrix Factorization Problems", 2017
59. Information Theory and Applications, "Low Rank Approximation with Entrywise L_1 -Norm Error", 2017
58. Asilomar Conference on Signals, Systems, and Computers, "Sketching as a Tool for Numerical Linear Algebra", 2016
57. Shanghai Tech Symposium on Information Science and Technology, "Sketching as a Tool for numerical Linear Algebra and Recent Developments", 2016
56. Simons Genomics Program Invited Talk, "Beating CountSketch for Heavy Hitters in Insertion Streams", 2016
55. SIAM Conference on Discrete Math, "Weighted Low Rank Approximation", 2016

54. Nexus of Information and Computation Theories Central Workshop, “Sketching as a Tool for Numerical Linear Algebra”, 2016
53. Institut Henri Poincare Inference Problems Theme, “New Algorithms for Heavy Hitters in Data Streams”, 2016
52. Institut Henri Poincare Fundamental Inequalities: techniques and applications, “Turnstile Streaming Algorithms Might as Well be Linear Sketches”, 2016
51. Alan Turing Institute Distributed Machine Learning and Optimization, “A Story of Principal Component Analysis in the Distributed Model”, 2015
50. IMA Resource Trade-Offs: Computation, Communication, and Information, “A Story of Principal Component Analysis in the Distributed Model”, 2016
49. Simons Computational Complexity of Low-Polynomial Time Problems, “Input Sparsity and Hardness for Linear Algebra Problems”, 2015
48. Johns Hopkins Sublinear Algorithms Workshop, “Beating CountSketch for Heavy Hitters in Insertion Streams”, 2016
47. NIPS BigNeuro Workshop, “Sketching as a Tool for Numerical Linear Algebra, and Recent Developments”, 2015
46. Shonan Recent Advances in Randomized Numerical Linear Algebra, “Optimal Principal Component Analysis in Distributed and Streaming Models”, 2016
45. Simons Real Time Decision Making Workshop, “Beating CountSketch for Heavy Hitters in Insertion Streams”, 2016
44. Information Theory and Applications, “Beating CountSketch for Heavy Hitters in Insertion Streams”, 2016
43. Information Theory and Applications, “The Sketching Complexity of Graph Cuts”, 2015.
42. Stanford Compression Forum, “Compressing Matrices”, 2015
41. Simons Big Data Reunion Workshop, “Turnstile Streaming Algorithms Might as Well Be Linear Sketches”, 2014
40. Distributed Machine Learning and Matrix Computations workshop at NIPS, “Principal Component Analysis and Higher Correlations for Distributed Data, 2014
39. Midwest Theory Day, “Principal Component Analysis and Higher Correlations for Distributed Data”, 2014

38. Simons Fast Algorithms via Spectral Methods Workshop, “Sketching as a Tool for Numerical Linear Algebra”, 2014
37. CERI/DIMACS Workshop on Streaming Graph Algorithms, “Numerical Linear Algebra in the Streaming Model”, 2014
36. Shonan Workshop for Large Scale Graphs, “The Sketching Complexity of Graph Cuts”, 2014
35. MIT/MSR Theory Day, “Turnstile Streaming Algorithms Might as Well be Linear Sketches”, 2014
34. Bertinoro Workshop on Sublinear Algorithms, “Turnstile Streaming Algorithms Might as Well be Linear Sketches”, 2014
33. Presburger Award Speech at ICALP, “Data Streams and Applications in Computer Science”, 2014
32. Berkeley Workshop on Algorithms for Modern Massive Data Sets, “Optimal CUR Matrix Decompositions”, 2014
31. UC San Diego Information Theory and Applications Workshop, “Sketching as a Tool for Numerical Linear Algebra”, 2014
30. Dagstuhl Workshop on Computational Complexity of Discrete Problems, “Turnstile Streaming Algorithms Might as Well be Linear Sketches”, 2014
29. Survey Talk at Banff Workshop on Communication Complexity and Applications, “Lower Bounds for Data Streams”, 2014
28. NIPS Workshop on Large Scale Matrix Analysis and Inference, “Sketching as a Tool for Numerical Linear Algebra”, 2013
27. Berkeley Simons Workshop on Succinct Data Representations and Applications, “How Robust are Linear Sketches to Adaptive Inputs?”, 2013
26. Aggregate Knowledge Conference on Data Streams, “Sketching as a Tool for Numerical Linear Algebra”, 2013
25. Coding, Complexity, and Sparsity Workshop at University of Michigan, “How Robust are Linear Sketches to Adaptive Inputs?”, 2013
24. Hong Kong Workshop on Theoretical Aspects of Big Data, “Low Rank Approximation and Regression in Input Sparsity Time”, 2013
23. STOC Workshop on Information Complexity and Applications, “Applications of Information

Complexity II”, 2013

22. Pacific Northwest Theory Day, “Low Rank Approximation and Regression in Input Sparsity Time”, 2013

21. Survey talk at Institute for Mathematics and its Applications (IMA) Workshop on Group Testing Designs, Algorithms, and Applications to Biology, “A Survey of Norm Estimation with Applications”, 2012

20. FOCS Workshop on Randomized Numerical Linear Algebra: Theory and Practice, “Low Rank Approximation and Regression in Input Sparsity Time”, 2012

19. Coding, Complexity, and Sparsity Workshop at University of Michigan, “Low Rank Approximation and Regression in Input Sparsity Time”, 2012

18. Workshop on Streaming Algorithms at University of Dortmund, “Low Rank Approximation and Regression in Input Sparsity Time”, 2012

17. Stanford Workshop on Algorithms for Modern Massive Data Sets, “Low Rank Approximation and Regression in Input Sparsity Time”, 2012

16. MSRI Workshop on Quantitative Geometry in Computer Science, “Tight Bounds for Distributed Functional Monitoring”, 2011

15. Banff Workshop on Information Theory and Statistics for Large Alphabets, “Sublinear Optimization for Machine Learning”, 2011

14. Survey talk at Bertinoro Workshop on Sublinear Algorithms, “A Survey of Norm Estimation with Applications”, 2011

13. Keynote Speaker at ACM SIGACT Chapter at Fudan University, “Efficient Sketches for Earthmover Distance and Cascaded Moments”, 2011

12. Coding, Complexity and Sparsity Workshop at University of Michigan, “ $(1+\epsilon)$ -Approximate Sparse Recovery”, 2011

11. Synergies in Lower Bounds Workshop at Aarhus University, “Lower Bounds in Communication for Streaming Problems Using Information-Theoretic Arguments”, 2011

10. Haifa Stringology Workshop, “Fast Moment Estimation in Data Streams in Optimal Space”, 2011

9. DIMACS Workshop on Network Data Streaming and Compressive Sensing, “An Optimal Algorithm for the Distinct Elements Problem”, 2010

8. Stanford Workshop on Algorithms for Modern Massive Data Sets, “Fast L_p Regression in

Data Streams”, 2010

7. Tel Aviv Stringology Workshop, “An Optimal Algorithm for the Distinct Elements Problem”, 2010

6. IITK Workshop on Algorithms for Data Streams, “1-Pass Relative Error L_p Sampling in a Data Stream”, 2010

5. Banff Workshop on Mathematics of String Spaces and Algorithmic Applications, “The Space Complexity of Cascaded Aggregates”, 2009

4. Dagstuhl Workshop on Sublinear Algorithms, “Revisiting Norm Estimation in Data Streams”, 2008

3. IPAM Workshop on Applications and Foundations of Cryptography and Computer Security on Private information Retrieval and Locally Decodable Codes, “New Lower Bounds for General Locally Decodable Codes”, 2006

2. DIMACS Workshop on Privacy-preserving Datamining, “Private Inference Control”, 2004

1. DIMACS Workshop on Discrete Metric Spaces and Algorithmic Applications, “Tight Lower Bounds for the Distinct Elements Problem”, 2003

Departmental Seminars and Colloquia

94. Stanford Information Systems Laboratory Colloquium, 2020 (upcoming)

93. UT Austin Theory Seminar, “Towards a Zero-One Law for Column Subset Selection”, 2020

92. National University of Singapore, “Towards a Zero-One Law for Column Subset Selection, 2020

91. Rice University, 2019, “Toward a Zero-One Law for Column Subset Selection”, 2019

90. MSR Redmond, 2019, “Towards a Zero-One Law for Column Subset Selection”, 2019

89. Purdue University, “Strong Coresets for k -Median and Subspace Approximation, Goodbye Dimensions”, 2019

88. Peking University, “Tight Bounds for L_1 Oblivious Subspace Embeddings”, 2019

87. Nanyang Technological Institute, “Tight Bounds for L_1 Oblivious Subspace Embeddings”, 2019

86. Northeast Normal University, “Relative Error Tensor Low Rank Approximation”, 2019

85. Carnegie Mellon University, “Strong Coresets for k-Median and Subspace Approximation, Goodbye Dimension”, 2019
84. Google Pittsburgh, “Sketching as a Tool for Numerical Linear Algebra”, 2019
83. University of Pennsylvania, “Sketching as a Tool for Numerical Linear Algebra”, 2019
82. Princeton, “Strong Coresets for k-Median and Subspace Approximation, Goodbye Dimension”, 2019
81. University of Minnesota, “Strong Coresets for k-Median and Subspace Approximation, Goodbye Dimension”, 2019
80. EPFL, “Weighted Low Rank Approximation with Provable Guarantees”, 2019
79. Google, “Tutorial on Sketching for Linear Algebra”, 2018
78. Lawrence Livermore Labs, “Relative Error Tensor Low Rank Approximation”, 2018
77. Tsinghua, “Relative Error Tensor Low Rank Approximation”, 2018
76. Chinese Academy of Sciences, “Parameterized Complexity of Matrix Factorization Problems”, 2018
76. Northeast Normal University, “Sketching as a Tool for Numerical Linear Algebra”, 2018
75. Chinese Academy of Sciences, “Sketching as a Tool for Numerical Linear Algebra”, 2018
74. University of Pennsylvania, “Relative Error Tensor Low Rank Approximation”, 2018
73. Indian Institute of Science, “The Parameterized Complexity of Matrix Factorization Problems”, 2017
72. CMU theory seminar, “Relative Error Tensor Low Rank Approximation”, 2017
71. MIT TOC Colloquium, “Relative Error Tensor Low Rank Approximation”, 2017
70. Google Zurich, "Sketching as a Tool for Numerical Linear Algebra", 2017
69. TU Dortmund, "Sketching as a Tool for Geometry Problems", 2017
68. Max Planck Institute, "Parameterized Complexity of Matrix Factorization Problems", 2017
67. Nanyang Technological University, “Sketching as a Tool for Numerical Linear Algebra”, 2017

66. Max Planck Institute, “Sketching as a Tool for Numerical Linear Algebra”, 2017
65. Google, “Low Rank Approximation with Entrywise L1-Norm Error” 2016
64. University of Washington at St. Louis, “Sketching as a Tool for Numerical Linear Algebra”, 2016
63. Alan Turing Institute London special masterclass, “Sketching as a Tool for Numerical Linear Algebra”, 2016
62. Comcast Labs Washington DC, “Sketching as a Tool for Numerical Linear Algebra”, 2016
61. Center for Communications Research (CCR), “Sketching as a Tool for Numerical Linear Algebra”, 2016
60. Chinese Academy of Sciences, “Beating CountSketch for Heavy Hitters in Insertion Streams”, 2016
60. University of Washington long board talk, “An Optimal Algorithm for Finding L2 Heavy Hitters”, 2016
59. University of Washington theory seminar, “Beating CountSketch for Heavy Hitters in Insertion Streams”, 2016
58. EPFL, “Beating CountSketch for Heavy Hitters in Insertion Streams”, 2016
57. University of Minnesota Digital Technology Center, “Sketching as a Tool for Numerical Linear Algebra”, 2016
56. IIT Madras, “Sketching as a Tool for Numerical Linear Algebra”, 2016
55. Georgia Tech, “An Optimal Algorithm for Finding L2 Heavy Hitters”, 2016
54. Georgia Tech, “Input Sparsity and Hardness for Robust Subspace Approximation”, 2015
53. Johns Hopkins University, “Sketching for M-Estimators: A Unified Approach to Robust Regression”, 2015
52. Rensselaer Polytechnic Institute, “Sketching as a Tool for Numerical Linear Algebra”, 2014
51. Princeton University, “The Sketching Complexity of Graph Cuts”, 2014
50. Duke University, “Turnstile Streaming Algorithms Might as Well be Linear Sketches”, 2014
49. University of Michigan, “The Sketching Complexity of Graph Cuts”, 2014

48. Microsoft Research Redmond, “Turnstile Streaming Algorithms Might as Well be Linear Sketches”, 2014
47. Chinese Academy of Sciences, “Data Stream Lower Bounds: a Tutorial”, 2014
46. National Institute of Informatics, Tokyo, “Turnstile Streaming Algorithms Might as Well be Linear Sketches”, 2014
45. TCS+ online seminar in theoretical computer science, “Turnstile Streaming Algorithms Might as Well be Linear Sketches”, 2014
44. University of British Columbia, “Sketching as a Tool for Numerical Linear Algebra”, 2014
43. Brown University, “Sketching as a Tool for Numerical Linear Algebra”, 2013
42. Purdue University, “Sketching as a Tool for Numerical Linear Algebra”, 2013
41. University of Indiana, “Sketching as a Tool for Numerical Linear Algebra”, 2013
40. Microsoft Silicon Valley, “Low Rank Approximation and Regression in Input Sparsity Time”, 2013
39. Georgia Tech, “Low Rank Approximation and Regression in Input Sparsity Time”, 2013
38. MIT, “Low Rank Approximation and Regression in Input Sparsity Time”, 2013
37. Stanford, “Low Rank Approximation and Regression in Input Sparsity Time”, 2013
36. Chinese Academy of Sciences, “Sketching Techniques for Numerical Linear Algebra”, 2012
35. Chinese Academy of Sciences, “Low Rank Approximation and Regression in Input Sparsity Time”, 2012
34. Tsinghua, “Low Rank Approximation and Regression in Input Sparsity Time”, 2012
33. Stanford, “Tight Bounds for Distributed Functional Monitoring”, 2012
32. Tokyo Institute of Technology, “Sampling and Sketching Techniques for Machine Learning”, 2012
31. University of Tokyo, “Additive Spanners”, 2012
30. IBM Almaden, “Tight Bounds for Distributed Functional Monitoring”, 2012
29. IBM Tokyo, “Sketching Techniques for Numerical Linear Algebra”, 2012

28. Iowa State University, "An Optimal Algorithm for the Distinct Elements Problem", 2011
27. Zhejiang University, "Sublinear Optimization for Machine Learning", 2011
26. Tsinghua, "Efficient Sketches for Earthmover Distance and Cascaded Moments", 2011
25. Google Tel Aviv, "Fast Robust Regression and Hyperplane Fitting", 2011
24. Microsoft Asia, "Sublinear Optimzation for Machine Learning", 2011
23. Technion, "Near-Optimal Private Approximation Protocols via a Black Box Transformation", 2011
22. Technion, "Subspace Embeddings for the L1-norm with Applications", 2011
21. Georgia Tech, "Subspace Embeddings for the L1-norm with Applications", 2011
20. Princeton University, "Subspace Embeddings for the L1-norm with Applications", 2011
19. IBM Almaden, "Efficient Sketches for Earth-Mover Distance, with Applications", 2010
18. University of Dortmund, "Efficient Sketches for Earth-Mover Distance, with Applications", 2010
17. University of Utah, "Fast Robust Regression in Data Streams", 2010
16. Universite Paris-Sud, "Fast Robust Regression in Data Streams", 2010
15. University of Dortmund, "Numerical Linear Algebra in a Data Stream", 2009
14. University of Bonn, "Transitive-Closure Spanners", 2008
13. Chinese University of Hong Kong, "Private and Efficient Distance Approximation", 2007
12. Toyota Institute of Technology, "Private and Efficient Distance Approximation", 2007
11. University of Southern California", "Private and Efficient Distance Approximation", 2007
10. IBM Almaden, "Private and Efficient Distance Approximation", 2007
9. Microsoft Research Silicon Valley, "Private and Efficient Distance Approximation", 2007
8. Stevens University, "Polylogarithmic Private Approximations and Efficient Matching", 2006
7. Tsinghua, "Explicit Exclusive Set Systems with Applications", 2006

6. City University of Hong Kong, "Lower Bounds for Additive Spanners, Emulators, and More", 2006
5. Chinese University of Hong Kong, "Lower Bounds for Additive Spanners, Emulators, and More", 2006
4. MIT, "Optimal Approximations of the Frequency Moments", 2005
3. Tsinghua, "Polylogarithmic Private Approximations and Efficient Matching", 2005
2. MIT, "Polylogarithmic Private Approximations and Efficient Matching", 2005
1. MIT, "Cryptography in an Unbounded Computational Model", 2002