BIOGRAPHICAL SKETCH AND PROFESSIONAL ACTIVITIES
RENSSELAER POLYTECHNIC INSTITUTE

February 5, 2026

I. Personal Data

Name:	Alex Gittens
Current Rank:	Associate Professor
Department:	Computer Science
School:	School of Science

Year and rank of first academic appointment at Rensselaer (e.g. instructor, assistant, associate, or full professor). Give dates and rank for subsequent promotions.

January 2017 - June 2025	Assistant Professor
July 2025 - present	Associate Professor

Date of Birth: November 1983

Educational Preparation

(1) Baccalaureate and graduate degree(s), institution, date

Ph.D.	(Applied and Comp. Mathematics)	CalTech, Pasadena, CA	2013
B.S.	(Mathematics)	University of Houston, Houston, Texas	2006
B.S	(Electrical Engineering)	University of Houston, Houston, Texas	2006

II. Professional Experience

(Give postdoctoral, teaching, industrial, governmental, and private practice experience prior to joining Rensselaer, giving position, employer, and dates.)

Postdoctoral Researcher	International Computer Science Institute, Berkeley, CA	April 2015–December 2016
Postdoctoral Researcher	eBay Research Laboratory San Jose, CA	June 2013–February 2015
Intern	IBM Research, Yorktown Heights, NY	Summer 2012

III. Teaching

A. Courses

(List the number and title of each course taught and approximate number of students in each course; Include undergraduate project supervision and courses that you have supervised with the number of instructors and students involved.)

Term	Course	Title	Instructor/course rating	Enrollment
2017 Spring	CSCI 6971/4971	Large Scale ML*	NA†/NA	5
2017 Fall	CSCI 6220/4030	Randomized Algorithms	(4.25/4.88)/(4.25/4.25)	14
2018 Spring	CSCI 6971/4971	Large Scale ML	(4.83/5.0)/NA	7
2018 Fall	CSCI 6220/4030	Randomized Algorithms	(4.5/4.5)/(4.0/3.5)	18
2019 Spring	CSCI 6971/4971	Large Scale ML	(4.25/4.67)/NA	7
2019 Fall	CSCI 6220/4030	Randomized Algorithms	(4.5/4.5)/(4.5/4.17)	18
2020 Spring	CSCI 2200	Foundations of CS	(3.67/3.13)$^\sharp$	222
2020 Fall	CSCI 6971/4971	ML and Optimization	(4/4.33)/(4.26/4.11)$^\sharp$	14/36
2021 Spring	CSCI 2200	Foundations of CS	(3.07/3.01)$^\sharp$	160
2021 Fall	CSCI 6971/4971	ML and Optimization	(4.57/4.57)/(4.45/4.55)	7/24
2022 Spring	CSCI 6968/4968	Trustworthy ML	(4.71/5.0)/(3.86/5.0)	10/3
2023 Spring	CSCI 2200	Foundations of CS	(2.59/2.50)	210
2023 Spring	CSCI 6968/4968	ML and Optimization	(4.25/4.4)/(4.13/4.14)	20/27
2023 Fall	CSCI 6968/4968	ML and Optimization	(3.29/3.29)/(3.5/3.5)	16/18
2024 Spring	CSCI 6968/4968	Trustworthy ML	(4.67/4.83)/(4.0/3.0)	7/2
2024 Fall	CSCI 4140/6962	ML and Optimization	(NA/NA)/(4.71/4.29)	13/16
2024 Fall	CSCI 6969	Graduate Skills	4.36/3.78	53
2025 Fall	CSCI 4180/6962	Trustworthy ML	(4/4)/(4/4.2)	2/12
2025 Fall	CSCI 6969	Graduate Skills	3.67/2.63	36

*The full name of the course is "Large Scale Matrix Computation and Machine Learning."

† NA indicates not available due to class size, or respondence rate.

$^\sharp$ This course was taught in the online modality during the pandemic.

Undergraduate Research Projects:

- Ashley Woodson, *Distribution Shift in Knowledge Graphs*, Summer 2024.
- Arvind Rathnashyam, *LLM Alignment via Corrector Models*, Summer 2024– Fall 2024.
- Arvind Rathnashyam, *Subquantile Minimization for Robust Kernel Learning*, Fall 2023.
- Hongcheng Zeng, *Robust ML via Subquantile Optimization*, Fall 2022.
- Hongcheng Zeng, *Anomaly Detection in Knowledge Graphs*, Spring 2022.
- Tim Budding, *Group Fair Federated Learning via Doubly-Stochastic Kernel-Target Alignment*, Fall 2021–Spring 2022.
- Jesse Ellin, *Fair and Sparse Generalized Additive Models*, Fall 2021.
- Gabriel Orlanski, *Improving Code Generation using Side information*, Spring and Summer 2021.
- Dylan Walker, *A generative approach to graph anomaly detection*, Spring and Summer 2021.
- Chris Jerrett, *Randomized First and Second-Order Tensor Low-Rank Approximation Algorithms*, Fall 2020.

- Tim Saia, *CUR for Matrix Completion*, Fall 2017.

B. **Student Thesis Supervision** (List student's name, title of thesis, and year completed.)

a. **Bachelors**
None

b. **Masters**

1. **Ryan Kaplan**, CSCI, master. Graduated Spring 2023.
2. **Jesse Ellin**, CSCI, master. Graduated Spring 2023.
3. **Kevin Kim**, CSCI, master. Graduated Fall 2021.
4. **Ngoc Truong**, CSCI, master. Graduated Spring 2021.
5. **Rufeng Ma**, CSCI, master. Graduated Spring 2021.

c. **Doctoral**

Ph.D. Students Advised:

Graduated

1. **Kareem Aggour**, CS. Co-advised with Bulent Yener. Graduated 2019. Thesis entitled "Intelligent and Scalable Canonical Polyadic Decomposition." Funding: GE, Inc.
2. **Dong Hu**, CS. Graduated 2024. Thesis entitled "Scalable Cost-Efficient Techniques for Machine Learning via Sketching." Funding: IBM AI Horizons Network Fellowship.
3. **Lilian Ngweta**, CS, Graduated 2025. Thesis entitled "Efficient Post-Processing Techniques for Foundation Models." Funding: IBM AI Research Collaboration.
4. **Sharmishtha Dutta**, CS, Graduated 2025. Thesis entitled "Curation and Completion of Knowledge Graphs." Funding: IBM AI Research Collaboration.

In Progress

1. **Huzaifa Arif**, ECSE, entered Ph.D. program Spring 2021. Graduation expected Spring 2026.
2. **Christopher Okonkwo**, CSCI, entered Ph.D. program Spring 2024. Graduation expected Fall 2027. Funding: Air Force Research Laboratory Fellowship.
3. **A F M Saif**, ECSE, entered Ph.D. program Fall 2023. Started working with me Fall 2025. Graduation expected Fall 2028.

Ph.D. Thesis Committee Memberships:

1. **Maksim Kholiavchenko**. CS. Advisor: Charles Stewart. Graduated Fall 2025.
2. **Arpan Mukherjee**. ECSE. Advisor: Ali Tajer. Graduated Summer 2024.
3. **Glory Justin**. ECSE. Advisor: Santiago Paternain. Graduated Spring 2024.
4. **Timothy Castiglia**. CS. Advisor: Stacy Patterson. Graduated Spring 2023.

5. **Ao Liu**. CS. Advisor: Lirong Xia. Graduated Spring 2023
6. **Farhad Mohsin**. CS. Advisor: Lirong Xia. Graduated Spring 2023.
7. **Mauricio Gruppi**. CS. Advisor: Sibel Adali. Graduated Fall 2022.
8. **Zichong Li**. Math. Advisor: Yangyang Xu. Graduated Fall 2022.
9. **Jun Wang**. CS. Advisor: Lirong Xia. Graduated Fall 2022.
10. **Zichong Li**. Math. Advisor: Yangyang Xu. Graduated Summer 2022.
11. **Daniel Park**. CS. Advisor: Bulent Yener. Graduated Spring 2021.
12. **Yu Chen**. CS. Advisor: Mohammed J. Zaki. Graduated Fall 2020.
13. **Brayden Hollis**, CS. Advisor: Stacy Patterson. Graduated Spring 2018.
14. **Nidhi Rastogi**, CS. Advisor: James Hendler. Graduated Spring 2018.

Masters Thesis Committee Memberships:

1. **Alex Mankowski**. CS. Advisor: Charles Stewart. Graduated Fall 2022.
2. **Maksim Kholiavchenko**. CS. Advisor: Charles Stewart. Graduated Spring 2022.
3. **Zhenhan Huang**. CS. Advisor: Jianxi Gao. Graduated Spring 2022.
4. **Shaunak Basu**. CS. Advisor: Malik Magdon-Ismail. Graduated Fall 2020.
5. **Eric Pianin**. CS. Advisor: Charles Stewart. Graduated Fall 2020.
6. **Andrew Batbouta**. CS. Advisor: Charles Stewart. Graduated Fall 2017.
7. **Nicholas Brennan**. CS. Advisor: Charles Stewart. Graduated Fall 2017.
8. **Dylan Elliott**. CS. Advisor: Mohammed Zaki. Graduated Fall 2017.

C. Course and Curriculum Development

CSCI 6971/4971, Large Scale Matrix Computation and Machine Learning

Spring 2017. I developed this new course for graduate students based on recent trends in the application of randomized dimensionality reduction techniques in numerical linear algebra and large-scale machine learning. The course provides an introduction to the concepts of large-scale stochatic optimization and matrix sketching, and highlights their applications to core computational primitives including the solution of regularized empirical risk minimization problems and the construction of low-rank approximation. It exposes students to the ways in which these fast optimization and linear algebraic primitives can then be exploited to trade-off between accuracy and training time when solving large-scale machine learning problems; kernel-based approaches to nonlinear clustering and regression are considered in detail. Students are expected to be familiar with the basics of computational linear algebra, probability, and algorithms. The course is theoretical in nature, with the goal being that the graduate students will be equipped to understand the proof of guarantees for randomized algorithms in the research literature, and the homeworks test and reinforce their facility with using these algorithms in practice. Near the end of the semester, the students are responsible for picking a recent paper introducing

an algorithm which uses the concepts covered in class, reproducing and extending the experimental verification of the algorithm's performance, commenting on the strength and weakness of the supplied theory, and presenting their results before the class. Grading is based on performance on the homeworks, in class pop-quizzes, and the project.

CSCI 6220/4030, Randomized Algorithms Fall 2017. I developed this randomized algorithms course based in part on the lecture notes of Yale University's CPSC 469/569, Jeff Erickson's (of UIUC) randomized algorithms lecture notes, and the Mitzenmacher–Upfal text. Compared to our previous Randomized Algorithms course, this course covered more in-depth the analysis of the performance of randomized algorithms. It covers the standard tools for the analysis of i.i.d. sums— Markov's and Chebyshev's inequalities, and Chernoff bounds— and applies these to the development and analysis of several classical randomized algorithms. Some of the classical topics covered include randomized selection, reservoir sampling, randomized routing, randomized data structures, markov chains, the probabilistic method, and applications of coupon collecting. The course finishes with a selection of modern randomized algorithms motivated by data science applications. Near the end of the semester, students pick a recent paper covering a randomized algorithm of their choice, reproduce and extend the experimental validation of the algorithm's performance, comment on the strength and applicability of any performance guarantees, and present their results before the class. Grading is based on performance on the homeworks, in class pop quizzes, and the project.

CSCI 6971/4140, Machine Learning and Optimization Fall 2020. I developed this new course for graduate students to introduce them to the foundations of the optimization and probability theory necessary to understand and use the optimization algorithms used in modern machine learning and scientific computing applications, as well as to introduce them to the basics of modern deep learning, from the first-order algorithms widely used therein to the common neural network architectures and regularization techniques. The course is a mix of theory and hands-on practice, with the first half of the course focusing on probability, convex optimization theory and algorithms, and the formulation of machine learning problems as convex optimization problems; and the second half focusing on the extension of these ideas to the nonconvex setting of machine learning with neural networks. In both parts, the homeworks consist of a mix of theoretical investigations and practical hands-on coding, using Numpy and the scikit-learn and pytorch machine learning frameworks to implement algorithms, analyze data sets, and train and evaluate deep neural networks. Near the end of the semester, the students are responsible for picking a recent paper introducing an algorithm which uses concepts covered in class, reproducing and extending the experimental verification of the algorithm's performance, commenting on the strength and weakness of the supplied theory, and presenting their results before the class. Grading is based on performance on the homeworks, weekly participation exercises designed to gauge student comprehension, and the project. The course was well-received, and I plan to teach it annually. The undergraduate version of the course entered the catalog in Fall 2024.

CSCI 6968/4180, Trustworthy ML Spring 2022. I co-developed and taught this course with Ali Tajer of the ECSE department. The course introduces students to

the large body of recent work addressing several desiderata that are of increasing importance as machine learning penetrates into more socially consequential areas: adversarial robustness, fairness, privacy, and interpretability. It is a seminar course, with the instructors providing a general introduction to the questions, framing, and approaches taken for each theme, and the students reading and presenting papers that delve into more focused detail. The students are also responsible for extending or profiling the performance of an existing algorithm for trustworthy ML and giving an in-class presentation. Grading is based on in-class participation and the clarity, thoroughness, and clarity of the students' presentations. The undergraduate version of this course entered the catalog in Fall 2024.

IV. Publications

A. Books, Monographs, Recordings, Large Scale Musical or Video Works, Commissions

(Give title, co-authors, or collaborators, if any, publisher or commissioners, and date. State if a contributing author to an edited compilation.)

1. Huzaifa Arif, Alex Gittens, and Pin-Yu Chen. Utility-Privacy Tradeoff in Federated Learning. In Xiaxio Li, Ziyue Xu, Huazhu Fu, editors, Federated Learning for Medical Imaging: Principles, Algorithms and Applications, chapter 9. Academic Press, first edition, 2024.

B. Journal Articles

(Give title, co-authors, if any, journal, volume, issue, date, paging. List spin-offs (i.e., publications essentially on same piece of work) from a major work under the major item. Provide copies for at least the past six years of all journal articles, abstracts, and book reviews if they are not already with the department file. This list should also contain articles accepted but not yet in print and those submitted and not yet reviewed.)

1. **In refereed journals and conferences (articles which are reviewed by peers in the field prior to publication).**

   (a) **Major articles**
   Total citations: 2347 (Google Scholar)
   Citation numbers for top 5 papers: 497, 158, 143, 125, 101
   Citation numbers for top 5 journal papers: 158, 143, 101, 74, 35

   Journal Publications (accepted)

   1. Huzaifa Arif, Alex Gittens, Pin-Yu Chen, Keerthiram Murugesan. Group Fair Federated Learning via Stochastic Kernel Regularization. Transactions on Machine Learning Research. 2025.
   2. Deniz Koyuncu, Alex Gittens, Bulent Yener, Moti Yung. Adversarial Missingness Attacks on Causal Structure Learning. Just Accepted. ACM Transactions on Intelligent Systems and Technology. 2023.
   3. Alex Gittens, Bulent Yener, Moti Yung. An Adversarial Perspective on Accuracy, Robustness, Fairness, and Privacy: Multilateral-Tradeoffs in Trustworthy

ML. IEEE Access. Volume 10, 2022.

4. Shusen Wang, Alex Gittens, Michael W. Mahoney. Scalable K-Means Clustering with Nyström Approximation: Relative-Error Bounds. Journal of Machine Learning Research. Volume 19, 2019.

5. Bo Liu, Liping Jing, Jia Li, Jian Yu, Alex Gittens, Michael W. Mahoney. Group Collaborative Representation for Image Set Classification. International Journal of Computer Vision. Volume 127(2), 2019.

6. Alex Gittens, Kai Rothauge, et al. Alchemist: An Apache Spark to MPI Interface. Concurrency and Computation: Practice and Experience. 2019.

7. Shushen Wang, Alex Gittens, Michael W. Mahoney. Sketched Ridge Regression: Optimization Perspective, Statistical Perspective, and Model Averaging. Journal of Machine Learning Research. Volume 18, 218:1–50, 2018.

8. Alex Gittens, Michael W. Mahoney. Revisiting the Nyström Method for Improved Large-Scale Machine Learning. Journal of Machine Learning Research. Volume 17, 117:1–65, 2016.

9. Christos Boutsidis, Alex Gittens. Improved Matrix Algorithms Via the Subsampled Randomized Hadamard Transform. SIAM Journal on Matrix Analysis and Applications. Volume 34 Issue 3, 1301–1340, 2013.

10. Richard Y. Chen, Alex Gittens, Joel A. Tropp. The Masked Sample Covariance Estimator: An Analysis Using Matrix Concentration Inequalities. Journal of Information and Inference: A Journal of the IMA. Volume 1 Issue 1, 2–20, 2012.

Refereed Conference Proceedings (accepted)

1. Deniz Koyuncu, Alex Gittens, Bulent Yener, Moti Yung. Exploiting Missing Data Remediation Strategies using Adversarial Missingness Attacks. In Proceedings of **AAAI Conference on Artificial Intelligence** (AAAI), 2026.

2. Huzaifa Arif, Pin-Yu Chen, Alex Gittens, Keerthiram Murugesan, Payel Das. PEEL the layers and find yourself: Revisiting Inference-Time Data Leakage for Residual Neural Networks. IEEE Secure and Trustworthy Machine Learning Conference (SaTML), 2025.

3. Sharmishtha Dutta, Alex Gittens, Mohammed Zaki, Charu Aggarwal. Replacing Paths with Connection-Biased Attention for Knowledge Graph Completion. In Proceedings of **AAAI Conference on Artificial Intelligence** (AAAI), 2025.

4. Lilian Ngweta, Mayank Agarwal, Subha Maity, Alex Gittens, Yuekai Sun, Mikhail Yurochkin. Aligners: Decoupling LLMs and Alignment. **Findings of the Association for Computational Linguistics**: EMNLP 2024.

5. Lilian Ngweta, Subha Maity, Alex Gittens, Yuekai Sun, Mikhail Yurochkin. Simple Disentanglement of Style and Content in Visual Representations. In Proceedings of **International Conference on Machine Learning** (ICML), 2023.

6. Deniz Koyuncu, Alex Gittens, Bulent Yener, Moti Yung. Deception by Omission: Using Adversarial Missingness to Poison Causal Structure Learning. In Proceed-

ings of **International Conference on Knowledge Discovery and Machine Learning** (KDD), 2023.

7. Huzaifa Arif, Alex Gittens, Pin-Yu Chen. Reprogrammable-FL: Improving Utility-Privacy Tradeoff in Federated Learning via Model Reprogramming. In Proceedings of IEEE Conference on Secure and Trustworthy Machine Learning, 2023.

8. Nidhi Rastogi, Sharmishtha Dutta, Mohammed Zaki, Alex Gittens, Charu Aggarwal. TINKER: A framework for Open source Cyberthreat Intelligence. 2022. 21st International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom'22), 2022.

9. Anik Saha, Alex Gittens, Jian Ni, Oktie Hassanzadeh, Bulent Yener, Krishava Srinivas. SPOCK@ Causal News Corpus 2022: Cause-Effect-Signal Span Detection Using Span-Based and Sequence Tagging Models. In Proceedings of the 5th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE), 2022.

10. Dong Hu, Shashanka Ubaru, Alex Gittens, Kenneth Clarkson, Lior Horesh, Vasileios Kalantzis. Sparse Graph-based Sketching for Fast Numerical Linear Algebra. In Proceedings of **International Conference on Acoustics, Speech and Signal Processing** (ICASSP), 2021.

11. Alex Gittens, Kareem Aggour, Bulent Yener. Adaptive Sketching for Fast and Convergent Canonical Polyadic Decomposition. In Proceedings of **International Conference on Machine Learning**, 2020.

12. Dong Hu, Alex Gittens, Malik Magdon-Ismail. NoisyCUR: An algorithm for two-cost budgeted matrix completion. In Proceedings of **European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases** (ECML\PKDD), 2020.

13. Nidhi Rastogi, Sharmishtha Dutta, Mohammed J Zaki, Alex Gittens, Charu Aggarwal. MALOnt: An Ontology for Malware Threat Intelligence. In Proceedings of **KDD MLHat Workshop on Deployable Machine Learning for Security Defense**, 2020.

14. Kareem Aggour, Alex Gittens, Bulent Yener. Accelerating a Distributed CPD Algorithm for Large Dense, Skewed Tensors. In Proceedings of **IEEE Big Data**, 2018.

15. Chander Iyer, Christopher Carothers, Petros Drineas, Alex Gittens. Randomized Block Iterative Solvers for Scalable Low Rank Approximation on Dense Terascale Matrices with Clustered Spectral Gaps. 9th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems **ScalA18**, held in conjunction with **SC18**.

16. Alex Gittens, Dimitris Achlioptas, Michael W. Mahoney. Skip-gram - Zipf + Uniform. In Proceedings of **Annual Meeting of the Assocation for Computational Linguistics** (ACL), 2017.

17. Shusen Wang, Alex Gittens, Michael W. Mahoney. Sketched Ridge Regression: Optimization Perspective, Statistical Perspective, and Model Averaging. In Proceedings of **International Conference on Machine Learning** (ICML), 2017.

18. Stephen Tu, Shivaram Venkataraman, Alisha Wilson, Alex Gittens, Michael Jordan, Benjamin Recht. Breaking Locality Accelerates Block Gauss–Siedel. In Proceedings of **International Conference on Machine Learning** (ICML), 2017.

19. Alex Gittens, Aditya Devarakonda, Evan Racah, et al. Matrix Factorization at Scale: a Comparison of Scientific Data Analytics in Spark and C+ MPI Using Three Case Studies. In Proceedings of **IEEE Big Data 2016**.

20. Alex Gittens, Jey Kottalam, Jiyan Yang, et al. A Multi-Platform Evaluation of the Randomized CX Low-Rank Matrix Factorization in Spark. In Proceedings of **International Parallel and Distributed Processing Symposium (IPDPS) ParLearning Workshop 2016**.

21. Christos Boutsidis, Alex Gittens, Prabhanjan Kambadur. Spectral Clustering via the Power Method - Provably. In Proceedings of **International Conference on Machine Learning** (ICML), 2015.

22. Da Kuang, Alex Gittens, Raffay Hamid. Hardware Compliant Approximate Image Codes. In Proceedings of **Conference on Computer Vision and Pattern Recognition** (CVPR), 2015.

23. Raffay Hamid, Ying Xiao, Alex Gittens, Dennis DeCoste. Compact Random Feature Maps. In Proceedings of **International Conference on Machine Learning** (ICML), 2014.

Refereed Conference Proceedings (submitted and in preparation)

1. Arvind Rathnashyam, Alex Gittens. Iterative Thresholding for Non-Linear Learning in the Strong ε-Contamination Model. *In preparation.*

(b) **Abstracts, letters of correspondence, book reviews, etc.**

None

(c) **Patents.**

None

V. **Research Grants and Contracts**
(Give title of project, other senior investigators, starting and completion dates, amount of funding, sponsoring agency.)

A. **Proposals Approved and Funded**

- Holistic Alignment of Agentic LLM Systems via Lightweight System-Level Objectives. PI: Alex Gittens. $110K. 1/1/2026–12/31/2026. Funded by FCRC IBM collaboration.
- Correctors and Selectors: Building An Ecosystem for LLM Alignment. PI: Alex Gittens. $100K. 1/1/2024–12/31/2024. Funded by AIRC IBM collaboration.

- Numerical Nonlinear Analysis for Deep Learning. (PhD research fellowship). PI: Alex Gittens. $273K. 8/6/2020–5/31/2023. Funded by AI Horizons Scholars Program.
- Anomaly Detection on Knowledge Graphs. co-PIs: Alex Gittens, Mohammed Zaki. $200K. 1/1/22–12/31/22. Funded by AIRC IBM collaboration.
- AI Models for Curation of Threat Intelligence. co-PIs: Alex Gittens, Mohammed Zaki, Nidhi Rastogi. $350K. 8/8/19–12/31/2021. Funded by AIRC IBM collaboration.

B. **Proposals submitted and not funded with current status**

(a) **Proposals pending.**

- PI Alex Gittens, coPI Bulent Yener. "SaTC 2.0: RES: Adversarial Missingness". NSF CISE SaTC. $886K, 5/18/2026–5/7/2029.

(b) **Unfunded proposals.**

- PI Alex Gittens, coPI Bulent Yener. "SaTC: CORE: Small: Adversarial Missingness". NSF CISE SaTC. $590K, 1/1/2025–12/31/2028.
- PI Alex Gittens, coPI Bulent Yener. "XAI for CO2 activation catalyst discovery". NSF CISE IIS. $410K, 8/1/2023–7/31/2026.
- PIs Alex Gittens, Rahmi Ozisik, Bulent Yener. "DMREF: Adaptive Adhesives". NSF Division of Materials Research. $1.1M, 2023–2024.
- PIs Alex Gittens, Bulent Yener, Ali Tajer. "Physics-Driven Causal Learning". DOE. $1.2M, 2021–2024.
- PI Alex Gittens. "CAREER: Next-generation Randomized Algorithms for Nonconvex Data Analysis". NSF CAREER, $600K, 2022–2027.
- PI Alex Gittens, coPI Bulent Yener. "Trustworthy Machine Learning— Algorithms Establishing Trade-offs between Accuracy, Robustness, Fairness, and Privacy." NSF SaTC Core. $1.2M. 2021–2024.
- PI Alex Gittens, coPI Bulent Yener. "Reducing bias in multimodal biomedical informatics via imputation and synthetic data generation", NIH, $2.5M, 2021–2025.
- PI Alex Gittens. "CAREER: Next-generation Randomized Algorithms for Nonconvex Data Analysis Problems". NSF CAREER, $675K, 2021–2025.
- PI Alex Gittens. "Numerical Nonlinear Algebra for Deep Learning". AIRC IBM collaboration, $200,000. 2020–2021
- PIs Alex Gittens and Bulent Yener. "Interpretability and Change-Point Prediction in Graphical Neural Networks". AIRC IBM collaboration, $200,000. 2020-–2021
- PI Lirong Xia, coPI Alex Gittens, coPI Malik Magdon-Ismail, coPI Elliot Anshelevich, coPI Sibel Adali. "AI Institute: Trustworthy AI", NSF AI Institute, $20M, 2020–2025.
- PI Alex Gittens, coPI Bulent Yener, coPI Malik Magdon-Ismail, coPI Karen Kafadar "A Theoretical Framework for Solving Structured Missing Data Problems", NSF-TRIPODS, $500,000, 2018–2021.

- PI Alex Gittens, coPI Bulent Yener, coPI Malik Magdon-Ismail, coPI Karen Kafadar "A Theoretical Framework for Solving Structured Missing Data Problems", NIH, $1,500,000, 2019–2023.

C. **Briefly Describe Your Current Research Interests**

I am interested in exploiting the trade-offs between computational effort and accuracy to design algorithms for the efficient solution of large-scale linear-algebraic and machine learning problems. In particular, I am interested in:

- The design and analysis of scalable machine learning algorithms
- Nonlinear and multilinear applications of sketching techniques
- Sampling for low-rank tensor/matrix approximation and completion
- Efficient approaches to trustworthy machine learning

A more detailed description of my research interests and objectives is given in my research statement.

VI. **Editorship of Journals, Reviews of Manuscripts, Books, and Research Proposals** (Give organization of journals, significant items reviewed, dates.)

(a) **Journal Editorship**

- Transactions on Machine Learning Research, Action Editor, 2022

(b) **Research Proposal Reviews**

- NSF Panelist in 2018, 2025
- Reviewer for WWTF (Vienna)

(c) **Journal Paper Reviews**

- **Computer Science:** VLDB Journal; Journal of Machine Learning Research; IEEE Transactions on Pattern Analysis and Machine Intelligence; Information and Inference: A Journal of the IMA, IEEE Transactions on Knowledge and Data Engineering
- **Applied Mathematics:** Advances in Computational Mathematics; SIAM Journal on Matrix Analysis and Applications; SIAM Journal on Scientific Computing; IEEE Transactions on Information Theory.

(d) **Conference/Workshop Chair**

- Program committee co-chair for the Scalable Deep Learning Workshop at IPDPS in 2022 and 2023.

(e) **Program Committees**

- **2026** Area chair, *AAAI Conference on Artificial Intelligence* (AAAI)
- **2020–2022, 2024** Area chair, *Conference on Neural Information Processing Systems* (NeurIPS)
- **2021,2023–2025** Area chair, *International Conference on Machine Learning* (ICML)
- **2020–2026** Area chair, *International Conference on Representation Learning* (ICLR)
- **2021** Senior program committee member, *International Joint Conference on Artificial Intelligence*, (IJCAI)

- **2020–2021** Program committee member, *International Conference on Knowledge Discovery and Data Mining* (KDD)
- **2019–2020,2024** Program committee member, *International Joint Conference on Artificial Intelligence* (IJCAI)
- **2018,2020** Program committee member, *International Conference on Artificial Intelligence and Statistics (AISTATS)*
- **2017,2019–2021** Program committee member, *AAAI Conference on Artificial Intelligence* (AAAI)
- **2017–2020** Program committee member, *International Conference on Machine Learning* (ICML)
- **2019** Program committee member, *Internation Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*
- **2018** Program committee member, *International Conference on High Performance Computing, Data, and Analytics (HIPC)*
- **2017–2019** Program committee member, *Conference on Neural Information Processing Systems* (NeurIPS)
- **2017,2018** Program committee member, *ACM International Conference on Information and Knowledge Management* (CIKM)

(f) **Conference Paper Reviews**

- (**in addition to conferences I served as PC above**) IEEE Annual Symposium on Foundations of Computer Science (FOCS), International Colloquium on Automata, Languages, and Programming (ICALP)

VII. Service

A. Service to University

1. *University, school, and departmental committees and dates for each.*
   - CS Admissions Committee (Fall 2025-)
   - CS Graduate Curriculum Committee (Fall 2024-)
   - CS Awards Committee (Fall 2024-)
   - CS Diversity and Inclusion Committee (Fall 2020-Fall 2021)
   - CS Colloquium Co-Chair (Fall 2017-Fall 2024)
   - CS Faculty Search Committe (Fall 2022-Spring 2023)
   - CS Undergraduate Curriculum Committee (Spring 2020-Fall 2022)
   - School of Science Dean's Review Committee (Fall 2021)
   - AI-ML Faculty Search Committee (Fall 2018-Fall 2019)
   - Advisory Committee for the RPI–IBM AI Research Collaboration (Fall 2017)
2. *Undergraduate student academic advising (number and year).*

   2017 41 students
   2018 47 students
   2019 42 students
   2020 44 students
   2021 65 students
   2022 59 students
   2023 58 students
   2024 64 students
   2025 33 students

3. *Graduate student academic advising (number and year).*

2018	1 student
2019	3 students
2020	4 students
2021	3 students
2022	6 students
2023	6 students
2024	5 students
2025	4 students

B. **Professional Societies**
(Give memberships, positions held, dates.)

- Member, AMS (2013–present)
- Member, SIAM (2013–present)
- Member, IEEE (2022–present)

VIII. **Professional and Public Lectures**
(List invited and contributed papers and lectures, giving title, organization, and dates.)

(a) **Invited Talks**

1. **University of Michigan-Dearborn REU Site in Mathematical Analysis and Applications**, University of Michigan at Dearborn, Dearborn, Michigan. Title: *Trading off accuracy for reduced computation in scientific computing*, 2024.
2. **Workshop on Sparse Tensor Computations, University of Illinois Urbana-Champaign**, Discovery Partners Institute, Chicago, Illinois. Title: *Faster Structured Tensor Decompositions via Sketching*, 2023.
3. **Institute for Mathematics and Its Applications, University of Minnesota**, Data Science Seminar. Title: *Trading off accuracy for reduced computation in scientific computing*, 2023.
4. **Computational Applied Mathematics and Operations Research, Rice University**, Houston, TX. Department colloquium. Title: *Sketching to trade-off accuracy for reduced computation in scientific computing*, 2023.
5. **Institute for Computational and Experimental Research in Mathematics, Brown University**, Providence, RI. Workshop on Modern Applied and Computational Analysis. Title: *Reduced label complexity for tight linear regression,* 2023.
6. **Operations Research and Information Engineering, Cornell University**, Ithaca, NY. Departmental Colloquium. Title: *Reduced label complexity for tight linear regression,* 2023.
7. **SIAM Conference on Computational Science and Engineering**, Randomized Numerical Algorithms for Matrix and Tensor Analysis Minisymposium. Title: *Acceleration of Randomized Tensor Decompositions,* 2023.
8. **Center for Applied Mathematics, Cornell**, Applied Mathematics Colloquim. Title *Intelligent Randomized Algorithms for the Low CP-Rank Tensor Approximation Problem,* 2022.

9. **Institute for Mathematics and Its Applications, University of Minnesota**, Data Science Seminar. Title *Intelligent Randomized Algorithms for the Low CP-Rank Tensor Approximation Problem,* 2022.
10. **University of Massachusetts, Lowell**, Machine Learning Seminar. Title: *Adaptive Sketching for the Low-rank Tensor Approximation Problem,* 2019.
11. **DIMACS Workshop on Randomized Numerical Linear Algebra, Statistics, and Optimization.** Rutgers University, New Brunswick, New Jersey. Title: *Adaptive Sketching for the Low-rank Tensor Approximation Problem,* 2019.
12. **SIAM Annual Meeting**, Randomized Numerical Linear Algebra Minisymposium. Title: *Sketched Ridge Regression: Kernel and Overdetermined Problems,* 2018.
13. **University of California, Berkeley**, Scientific Computing and Matrix Computations Seminar Series, Title: *Terabyte-scale Linear Algebra: Spark vs MPI, and Spark with MPI,* 2016.
14. **Nii Shonan Meeting**, Shonan Village Center, Japan, Recent Advances in Randomized Linear Algebra, Title: *Low-Rank Factorizations at Scale,* 2016.
15. **Park City Mathematics Institute**, Mathematics of Data Summer Program, Title: *Low-Rank Factorizations at Scale,* 2016.
16. **International Computer Science Institute**, Lunch Talk Series, Title: *Compositional Properties of Nonlinear Embeddings,* 2016.
17. **University of Colorado, Boulder**, Applied Mathematics Colloqium, Title: *Why (Some) Nonlinear Embeddings Capture Compositionality Linearly,* 2015.
18. **University of Michigan**, Ann Arbor, MI, Geometric Functional Analysis Group Meeting, Title: *Estimation of Sparse Covariance Matrices,* 2011.

(b) **Contributed Talks**

1. **Symposium on Data Science and Statistics**, Reston, Virginia. Title: *Optimal Dimensionality Reduction for Non-Linear Clustering Via Nystrom Approximation,* 2018.
2. **Conference of the International Linear Algebra Society**, Providence, Rhode Island, Title: *Randomized Low-Rank Approximations, in Theory and Practice*, 2013.
3. **Workshop on Signal Processing with Adaptive Sparse Structured Representations (SPARS)**, Edinburgh, UK, Title: *Tail Bounds for All Eigenvalues of a Sum of Random Matrices,* 2011.

IX. Honors and Awards

(List names and dates)

None

X. Sabbatical leaves, off-campus study programs, foreign professional travel

(Give dates and topics.)

None

XI. **Other Activities**

(List other relevant activities such as consulting, (include name of company and days per year), expert witness, or significant activities not included in previous categories).

None

Signature:  Date: February 5, 2026