



**Digital Research
Alliance of Canada**



Protected when completed

This is a draft version only. Do not submit to any funding organization. Only the final version from the History page can be submitted.

Professor Margo Seltzer

Correspondence language: English

Sex: Female

Contact Information

The primary information is denoted by (*)



**Digital Research
Alliance of Canada**



Protected when completed

This is a draft version only. Do not submit to any funding organization. Only the final version from the History page can be submitted.

Professor Margo Seltzer

Degrees

Doctorate, Doctor of Philosophy, Computer Science, University of California, Berkeley
Degree Status: Completed

Bachelor's, Bachelor of Arts, Applied Mathematics, Harvard University
Degree Status: Completed

Recognitions

- 2023/5 ACM Athena Lecturer Award - 25,000
Association for Computing Machinery
Prize / Award
This award celebrates women researchers who have made fundamental contributions to Computer Science. Each year ACM honors a preeminent woman computer scientist as the Athena Lecturer. The recipient gives an invited talk at a major ACM conference of her choice.
- 2023/5 Bell Labs Prize Phase 2 Selection
Nokia Bell Labs
Prize / Award
Every year we award the Bell Labs Prize to an innovator that we believe will lay the foundations of the next technology revolution and in the process, manifestly change the way we live, work and communicate.
- 2023/5 Killam Teaching Award
University of British Columbia
Prize / Award
As one of six Killam institutions, UBC offers yearly awards from the Killam Endowment Fund to faculty and teaching assistants who demonstrate excellence in teaching.
- 2022/11 Spotlight for "Anomaly Detection in Multiplex Dynamic Networks: from Blockchain Security to Brain Disease Prediction."
NeurIPS
Honor
NeurIPS 2022 Workshop on Temporal Graph Learning
- 2022/11 Best paper award for "Tinkertoy: Build your own operating system for IoT devices."
Association for Computing Machinery
Prize / Award
ACM Embedded Systems Conference (2022)

- 2022/11 Finalist of Data Mining Best Paper Competition Award (student track), INFORMS 2022
NeurIPS
Prize / Award
For our work published as an Oral Presentation at NeurIPS 2022, "Exploring the Whole Rashomon Set of Sparse Decision Trees."
- 2022/11 ACM Systems Award - 35,000
Association for Computing Machinery
Prize / Award
Awarded to an institution or individual(s) recognized for developing a software system that has had a lasting influence, reflected in contributions to concepts, in commercial acceptance, or both.
- 2022/11 Oral presentation for "Exploring the Whole Rashomon Set of Sparse Decision Trees."
NeurIPS
Honor
NeurIPS 2022
- 2021/5 Positive Teaching Letter
University of British Columbia
Prize / Award
Faculty of Science
- 2021/5 Elected to the American Academy of Arts and Sciences
American Academy of Arts and Sciences
Honor
The American Academy of Arts and Sciences is one of the oldest learned societies in the United States. Academy members are world leaders in the arts and sciences, business, philanthropy, and public affairs.
- 2020/6 SIGMOD Systems Award
University of California, Berkeley
Prize / Award
to Berkeley DB
- 2020/5 Incredible Instructor Award
University of British Columbia
Prize / Award
departmental award in computer science
- 2019/7 USENIX Lifetime Achievement Award
USENIX
Prize / Award
Margo Seltzer received the 2019 award in recognition of her research into experimental file and storage systems, her development of new storage paradigms such as provenance, her software contributions, and her dedication to and steering of the USENIX community and its organization.
- 2019/2 Elected to the (US) National Academy of Engineering
(US) National Academy of Engineering
Prize / Award
Elected to the (US) National Academy of Engineering

2017/2 CRA-E Undergraduate Research Mentoring Award
 Computing Research Association
 Prize / Award
 The CRA-E Undergraduate Research Faculty Mentoring Award recognizes individual faculty members who have provided exceptional mentorship, undergraduate research experiences and, in parallel, guidance on admission and matriculation of these students to research-focused graduate programs in computing.

Employment

2018/9 Canada 150 Research Chair in Computer Systems and Cheriton Family Chair in
 Computer Science
 Computer Science, University of British Columbia
 Full-time, Professor
 Tenure Status: Tenure

2017/7 - 2018/12 Instructor
 Online Business Analytics Program, Harvard Business School
 Part-time, Term, Professor
 Tenure Status: Tenure
 Developed and delivered online course materials and weekly online classes.

2017/7 - 2018/8 Visiting Professor
 Computer Science, Harvard Business School
 Full-time
 Tenure Status: Tenure

2006/3 - 2018/8 Architect
 Oracle Corporation

2000/7 - 2018/8 Herchel Smith Professor of Computer Science
 School of Engineering and Applied Sciences, Computer Science, Harvard University
 Full-time, Professor
 Tenure Status: Tenure

2005/9 - 2010/9 Harvard College Professor
 Computer Science, Harvard University
 Full-time
 Tenure Status: Tenure

2002/9 - 2006/2 Associate Dean
 Computer Science and Engineering, Harvard University
 Full-time
 Tenure Status: Tenure
 Responsible for overall administration of Computer Science and Electrical Engineering including curricular planning, Jr. faculty recruiting and mentoring, departmental communication, and Industrial outreach

1996/6 - 2006/2 Chief Technical Officer
 Sleepycat Software

2000/7 - 2004/9 Gordon McKay Professor
 Computer Science, Harvard University
 Full-time
 Tenure Status: Tenure

1997/7 - 2000/6	Associate Professor Division of Engineering and Applied Sciences, Computer Science, Harvard University Full-time, Associate Professor Tenure Status: Tenure Track
1993/1 - 1997/6	Assistant Professor Division of Engineering and Applied Sciences, Computer Science, Harvard University Full-time, Assistant Professor Tenure Status: Tenure Track
1988/7 - 1992/6	Research Assistant University of California, Berkeley
1986/10 - 1987/4	Senior Engineer Kendall Square Research Corp Responsible for VLSI logic design and simulation of full custom proprietary RISC CPU. Participated in design of instruction set architecture and interchip communication protocols. Also responsible for tape out of Kendall Square's first chip including DRC, ERC, layout versus schematic verification and MEBES generation. Designed and managed the implementation of automated CAD environment.
1985/8 - 1986/9	Senior Programmer Stratus Computer Corp Member of the file system and transaction processing group. Designed and implemented new queueing mechanisms to provide increased concurrency in the operating system. Implemented performance enhancements to the transaction processing subsystem and performance metering facilities in the file system.
1983/1 - 1985/7	Senior Engineer Sequoia Systems, Inc

Research Funding History

Awarded [n=15]

2018/9 - 2025/8 Principal Investigator	Canada 150 Research Chair in Computer System, Research Chair Funding Sources: Natural Sciences and Engineering Research Council of Canada (NSERC) Canada 150 Research Chairs Total Funding - 7,000,000
2019/1 - 2024/12 Principal Investigator	VELOCITY: A New Architecture for OS Design and Implementation, Grant Funding Sources: Natural Sciences and Engineering Research Council of Canada (NSERC) Total Funding - 240,000
2022/1 - 2022/12 Principal Investigator	OS Isolation, Grant Funding Sources: ARM Total Funding - 62,000
2022/1 - 2022/12 Principal Investigator	Device Driver Synthesis, Grant Funding Sources: Huawei Total Funding - 145,000
2022/1 - 2022/12	Graph Storage and Analysis, Grant

Principal Investigator **Funding Sources:**
Huawei
Total Funding - 133,000

2022/1 - 2022/12 Provenance and ML, Grant
Principal Investigator **Funding Sources:**
Oracle
Total Funding - 64,500

2021/1 - 2021/12 ML Reproducibility, Grant
Principal Investigator **Funding Sources:**
Oracle
Total Funding - 62,523

2019/11 - 2021/12 Graph Storage and Analytics: From the Bottom Up, Grant
Principal Investigator **Funding Sources:**
Huawei
Total Funding - 131,876

2020/1 - 2020/12 VELOCITY: Bringing the Benefits of VLSI to System Software, Grant
Principal Investigator **Funding Sources:**
Canada Foundation for Innovation (CFI)
Total Funding - 453,594

2020/1 - 2020/12 VELOCITY: Bringing the Benefits of VLSI to System Software, Grant
Principal Investigator **Funding Sources:**
British Columbia Knowledge Development Fund
Total Funding - 453,594

2018/9 - 2020/8 Increasing Scientific Dataset Quality Through Reproducibility and Curation Tools and
Co-applicant Targeted Services in Dataverse Repositories, Grant
Funding Sources:
Sloan Foundation
Total Funding - 499,697

2014/8 - 2020/7 XPS: FULL: CCA: Collaborative Research: Automatically Scalable Computation, Grant
Principal Investigator **Funding Sources:**
National Science Foundation (USA)
Exploiting Parallelism and Scalability (XPS)
Total Funding - 525,000
National Science Foundation (USA)
Exploiting Parallelism and Scalability (XPS)
Total Funding - 115,000

2015/6 - 2020/5 CISE-Provenance : SI2-SSI: Collaborative Research: Bringing End-to-End Provenance to
Principal Investigator Scientists, Grant
Funding Sources:
National Science Foundation (USA)
Software Infrastructure for Sustained Innovation
Total Funding - 1,422,728

2015/11 - 2019/11 PRINCESS : Probabilistic Representation of Intent Commitments to Ensure Software
Principal Investigator Survival (PRINCESS), Contract
Funding Sources:
Defense Advanced Research Project Agency (The)

	BRASS Total Funding - 1,380,154
2018/1 - 2018/12 Co-investigator	Towards a FAIR Digital Ecosystem in the Cloud, Grant Funding Sources: National Institutes of Health (NIH) (USA) Total Funding - 647,221
Completed [n=5]	
2017/9 - 2018/8 Principal Investigator	New Approaches for Ranking in Machine Learning, Grant Funding Sources: Duke University (USA) Subcontract from IIS-1053407 Total Funding - 49,767
2015/1 - 2017/12 Principal Investigator	Citation++: Data citation, provenance, and documentation, Grant Funding Sources: National Science Foundation (USA) Total Funding - 300,000
2013/10 - 2017/9 Principal Investigator	CSR: Medium: Collaborative Research: Workload-Aware Storage Architectures for Optimal Performance and Energy Efficiency, Grant Funding Sources: National Science Foundation (USA) Computer Systems Research Total Funding - 306,077
2009/9 - 2013/9 Co-investigator	Analyzing Complex Healthcare Data to Determine Causality of Observed Drug Effects, Grant Funding Sources: National Institutes of Health (NIH) (USA) Total Funding - 213,491
2009/10 - 2012/9 Principal Investigator	Collaborative Research: Scalable Data Management Using Metadata and Provenance, Grant Funding Sources: National Science Foundation (USA) Computing and Communication Foundations Total Funding - 351,643

Student/Postdoctoral Supervision

Bachelor's [n=35]

2023/4 - 2023/8 Principal Supervisor	Sraavan Sridhar, University of British Columbia Present Position: Undergraduate Student
2023/4 - 2023/8 Principal Supervisor	Emily Chu, University of British Columbia Present Position: Undergraduate Student
2023/1 - 2023/8 Co-Supervisor	Ryan Mehri, University of British Columbia Present Position: Undergraduate Student

2023/1 - 2023/8 Principal Supervisor	Jennifer Wong, University of British Columbia Present Position: Undergraduate Student
2022/9 - 2023/4 Co-Supervisor	Bryce Wilson, University of British Columbia Present Position: Undergraduate Student
2022/9 - 2022/12 Principal Supervisor	Ryan Liu, University of British Columbia Present Position: Undergraduate Student
2022/4 - 2023/4 Principal Supervisor	Frederick Shpilevskiy, University of British Columbia Present Position: Undergraduate Student
2022/4 - 2022/8 Principal Supervisor	Portia Chang, University of British Columbia Present Position: Undergraduate Student
2022/4 - 2022/8 Principal Supervisor	David Bromley, University of British Columbia Present Position: Undergraduate Student
2022/4 - 2023/8 Principal Supervisor	Sepehr Nooraafshan, University of British Columbia Present Position: Undergraduate Student
2022/4 - 2022/8 Principal Supervisor	Jack Li, University of British Columbia Present Position: Undergraduate Student
2022/4 - 2022/12 Principal Supervisor	Haotian Gong, University of British Columbia Present Position: Undergraduate Student
2022/1 - 2022/8 Principal Supervisor	Petal Vitis, University of British Columbia Present Position: Undergraduate Student
2021/9 - 2022/4 Principal Supervisor	Richard Zhu, University of British Columbia Present Position: Undergraduate Student
2021/4 - 2023/8 Principal Supervisor	Ilias Karimalis, University of British Columbia Present Position: Undergraduate Student
2021/4 - 2021/8 Principal Supervisor	Jacques Chen, University of British Columbia Present Position: Undergraduate Student
2021/1 - 2023/4 Principal Supervisor	Maxwell (Shukan) Yang, University of British Columbia Present Position: Undergraduate Student
2021/1 - 2022/4 Principal Supervisor	Romina Mahinpei, University of British Columbia Present Position: Undergraduate Student
2020/5 - 2020/8 Principal Supervisor	Nichole Bouffard, University of British Columbia Present Position: Undergrad student
2020/5 - 2020/8 Principal Supervisor	Tianhang Cui, University of British Columbia Present Position: Undergrad student
2020/5 - 2020/8 Principal Supervisor	Akash Sivaram, ITT Delhi Present Position: Undergrad student
2020/5 - 2020/8 Principal Supervisor	Alexander Zheng, University of British Columbia Present Position: Undergrad student
2020/5 - 2020/8 Principal Supervisor	Wesley Ferguson, University of British Columbia Present Position: Undergrad student
2020/4 - 2020/12 Principal Supervisor	Jude Shamsi, University of British Columbia Present Position: Undergraduate Student
2019/9 - 2020/4 Principal Supervisor	Narun Raman (Completed) , Carleton College Present Position: unknown

2019/9 - 2020/5 Principal Supervisor	Erik Carlson, Carleton College Present Position: Undergrad student
2019/5 - 2019/8 Principal Supervisor	Ashish Nair (Completed) , IIT Delhi Present Position: unknown
2019/4 - 2019/8 Principal Supervisor	Rachel Yeo, University of British Columbia Present Position: Undergraduate Student
2019/4 - 2019/8 Principal Supervisor	Rui Zhang, University of British Columbia Present Position: Undergraduate Student
2019/4 - 2019/8 Principal Supervisor	Gwang Chul Kim, University of British Columbia Present Position: Undergraduate Student
2019/4 - 2023/8 Principal Supervisor	Hayden McTavish, University of British Columbia Present Position: Undergraduate Student
2019/4 - 2019/8 Principal Supervisor	Frank Yan, University of British Columbia Present Position: Undergraduate Student
2019/4 - 2019/8 Principal Supervisor	Prateek Agarwal, University of British Columbia Present Position: Undergraduate Student
2019/1 - 2020/8 Principal Supervisor	Natalie Ngan, University of British Columbia Present Position: Undergrad student
2018/1 - 2018/4 Principal Supervisor	Jeanette Johnson, University of British Columbia Present Position: unknown

Master's Thesis [n=14]

2023/9 - 2025/5 Principal Supervisor	Sadaf Sadeghian, University of British Columbia Present Position: MSc student
2021/9 - 2023/11 Principal Supervisor	Hadi Sinaee (In Progress) , University of British Columbia Present Position: MSc student
2021/9 - 2023/11 Principal Supervisor	Milad Rezaei (In Progress) , University of British Columbia Present Position: MSc student
2021/9 - 2023/11 Principal Supervisor	Zainab Wattoo (In Progress) , University of British Columbia Present Position: MSc student
2020/9 - 2023/11 Principal Supervisor	Alexander Trostanovsky (In Progress) , University of British Columbia Present Position: MSc student
2020/9 - 2023/11 Principal Supervisor	Ali Behrouz, University of British Columbia Present Position: MSc student
2019/9 - 2021/5 Co-Supervisor	Junfeng Xu (Completed) , UBC Present Position: unknown
2019/9 - 2020/4 Principal Supervisor	Marie Pauline Sauvart (Completed) , University of Bonn Present Position: unknown
2019/9 - 2020/5 Principal Supervisor	Marie Pauline Sauvart, University of Bonn Present Position: Master's student
2019/9 - 2021/5 Principal Supervisor	Joseph Wonsil, University of British Columbia Present Position: PhD student
2019/9 - 2021/5 Principal Supervisor	Bingyao Wang (In Progress) , UBC Student Degree Expected Date: 2021/5 Present Position: MS Student

2019/9 - 2021/5
Principal Supervisor Michael Kim (In Progress) , UBC
Student Degree Expected Date: 2021/5
Present Position: MS Student

2018/9 - 2020/5
Academic Advisor Zixuan Yin (Completed) , The University of British Columbia
Present Position: unknown

2018/5 - 2020/9
Co-Supervisor Christopher Chen (In Progress) , The University of British Columbia
Student Degree Expected Date: 2020/4
Present Position: MSc Student

Doctorate [n=25]

2022/9 - 2027/5
Co-Supervisor Mohammad Dashti (In Progress) , University of British Columbia
Present Position: PhD student

2021/9 - 2026/5
Principal Supervisor Shaurya Patel (In Progress) , University of British Columbia
Present Position: PhD student

2021/9 - 2026/5
Principal Supervisor Sid Agrawal (In Progress) , University of British Columbia
Present Position: PhD student

2021/9 - 2025/5
Co-Supervisor Joel Nider (In Progress) , University of British Columbia
Present Position: PhD student

2021/9 - 2026/5
Co-Supervisor Christopher Chen (In Progress) , University of British Columbia
Present Position: PhD student

2021/9 - 2026/5
Principal Supervisor Bingyao Wang (In Progress) , University of British Columbia
Present Position: PhD student

2021/1 - 2021/12
Principal Supervisor Patrick Colp (Completed) , University of British Columbia
Present Position: unknown

2020/9 - 2025/5
Principal Supervisor Puneet Mehrotra (In Progress) , University of British Columbia
Present Position: PhD student

2020/9 - 2025/5
Principal Supervisor Chudi Zhang (In Progress) , University of British Columbia
Present Position: PhD student

2019/9 - 2025/5
Principal Supervisor Joseph Wonsil (In Progress) , UBC
Student Degree Expected Date: 2025/5
Present Position: Ph.D. Student

2019/9 - 2020/5
Principal Supervisor Maryam Raiyat, University of Tehran
Present Position: PhD student

2018/11 - 2022/5
Co-Supervisor Surbhi Palande (In Progress) , UBC
Student Degree Expected Date: 2022/5
Present Position: Ph.D. Student

2018/9 - 2020/5
Principal Supervisor Tony Mason (In Progress) , Harvard School of Engineering and Applied Sciences
Present Position: PhD Student

2018/5 - 2020/9
Co-Supervisor Swati Goswami (In Progress) , Harvard School of Engineering and Applied Sciences
Present Position: PhD Student

2016/9 - 2020/5
Co-Supervisor Crystal Hu (In Progress) , Harvard School of Engineering and Applied Sciences
Present Position: PhD Student

2016/9 - 2020/5
Co-Supervisor Michael Han (In Progress) , Harvard School of Engineering and Applied Sciences
Present Position: PhD Student

2016/9 - 2022/5 Co-Supervisor	Jingmei Hu (Completed) , University of British Columbia Present Position: unknown
2016/9 - 2022/5 Co-Supervisor	Xueyuan Han (Completed) , University of British Columbia Present Position: unknown
2015/9 - 2020/5 Co-Supervisor	David Holland (In Progress) , Harvard School of Engineering and Applied Sciences Present Position: PhD Student
2014/9 - 2020/5 Principal Supervisor	Robert Bowden (In Progress) , Harvard School of Engineering and Applied Sciences Present Position: PhD Student
2011/9 - 2014/8 Principal Supervisor	Elaine Angelino (Completed) , Harvard School of Engineering and Applied Sciences Present Position: Independent
2010/9 - 2012/5 Academic Advisor	Jason Waterman, Harvard School of Engineering and Applied Sciences Present Position: Assistant Professor
2008/9 - 2017/8 Principal Supervisor	Daniel Margo (Completed) , Harvard School of Engineering and Applied Sciences Present Position: Member of the Technical Staff
2008/9 - 2015/3 Principal Supervisor	Peter Macko (Completed) , Harvard School of Engineering and Applied Sciences Present Position: Member of the Technical Staff
2004/9 - 2014/5 Principal Supervisor	Uri Braun (Completed) , Harvard School of Engineering and Applied Sciences Present Position: CEO

Post-doctorate [n=9]

2021/9 - 2022/11 Principal Supervisor	Amee Trivedi (Completed) , University of British Columbia Present Position: unknown
2021/7 - 2022/9 Principal Supervisor	Maryam Aliabadi (Completed) , University of British Columbia Present Position: unknown
2021/2 - 2023/7 Principal Supervisor	Arpan Gujarati (In Progress) , University of British Columbia Present Position: Post-doctoral research fellow
2020/12 - 2023/12 Principal Supervisor	Reto Acherman (In Progress) , University of British Columbia Present Position: Post-doctoral research fellow
2017/9 - 2020/7 Co-Supervisor	Berk Ustun, Harvard School of Engineering and Applied Sciences Present Position: Post doctoral scholar
2016/9 - 2019/9 Co-Supervisor	Ming Kawaguchi (In Progress) , Harvard School of Engineering and Applied Sciences Present Position: Post doctoral researcher
2016/9 - 2017/12 Principal Supervisor	Thomas Pasquier (Completed) , Harvard School of Engineering and Applied Sciences Present Position: Lecturer
2014/7 - 2016/6 Principal Supervisor	Jacob Whitehill (Completed) , Harvard School of Engineering and Applied Sciences Present Position: Professor
2010/10 - 2012/10 Principal Supervisor	Marc Chiarini, Harvard School of Engineering and Applied Sciences Present Position: Senior Performance Engineer

International Collaboration Activities

2018/9 - 2021/5	Advisor, United States of America I continue to advise four Ph.D. students at Harvard University.
-----------------	--

- 2018/1 - 2020/12 Researcher, United Kingdom
I collaborate with Robert Watson at Cambridge University and Thomas Pasquier at the University of Bristol in the area of data provenance.
- 2016/1 - 2019/12 Researcher, United States of America
I conduct research on interpretable machine learning with Cynthia Rudin at Duke University, students at Harvard University, and Elaine Angelino from Berkeley.
- 2017/10 - 2018/10 Organizer, Germany
With colleagues from New Zealand (David Ayers) and the UK (Jatindra Singh and Christopher Millard), I co-organized a Dagstuhl Workshop on accountable systems.

Presentations

1. (2023). Distinguished Lecture. Duke University, Durham, United States of America
Invited?: Yes, Keynote?: Yes
2. (2023). Keynote. CRA-WP Grad Cohort, United States of America
Invited?: Yes, Keynote?: Yes
3. (2022). Applications of Data Provenance. Oracle Seminar Series (virtual), Canada
Invited?: Yes, Keynote?: No
4. (2022). Keynote presentation. Programming Language Design and Implementation conference (PLDI), San Diego, United States of America
Invited?: Yes, Keynote?: No
5. (2022). Invited Talk. High Performance Transaction Processing Workshop, Asilomar, United States of America
Invited?: Yes, Keynote?: No
6. (2021). Caching: It's not just for Data. Distinguished Lecture: MPI-SWS (virtual), Canada
Invited?: Yes, Keynote?: Yes
7. (2021). An NVM Carol. Distinguished Lecture: ETH Zurich (virtual), Switzerland
Invited?: Yes, Keynote?: No
8. (2021). NVM: Bubble Memory all over Again?. Dagstuhl Seminar (virtual), Canada
Invited?: Yes, Keynote?: No
9. (2021). CS313: Before and After. UIUC Workshop on Pandemic Teaching (virtual), Canada
Invited?: Yes, Keynote?: No
10. (2021). Applications of Data Provenance. NetApp Seminar Series, Canada
Invited?: Yes, Keynote?: No
11. (2021). Building your own 3-legged Stool. CRA Graduate Cohort (virtual), Canada
Invited?: Yes, Keynote?: Yes
12. (2021). When Databases met UNIX: A Love Affair in Five Acts. University of Sydney, John Lions Symposium Keynote Speaker (virtual), Canada
Invited?: Yes, Keynote?: Yes
13. (2020). End to End Provenance. NSF SSI PI meeting, Seattle, United States of America
Invited?: Yes, Keynote?: No
14. (2020). The Fine Line between Bold and Fringe Lunatic. Usenix Annual Technical Conference (USENIX ATC), Virtual, United States of America
Main Audience: Researcher
Invited?: Yes, Keynote?: Yes

15. (2019). Distinguished Lecture. An NVM Carol. Carleton College, Minnesota, United States of America
Invited?: Yes, Keynote?: No
16. (2019). An NVM Carol: Visions of NVM Past, Present, and Future. Alumni/Industry Lecture, California, United States of America
Invited?: Yes, Keynote?: No
17. (2019). Invited Lecture. Automatically Scalable Computation. EPFL, Lausanne, Switzerland
Invited?: Yes, Keynote?: No
18. (2019). Invited speaker. Systems Research - Construed Broadly. Bristol University Cybersecurity Colloquium, Bristol, United Kingdom
Invited?: Yes, Keynote?: No
19. (2019). Distinguished Lecture. Systems Research - Construed Broadly. University of Waterloo, Ontario, Canada
Invited?: Yes, Keynote?: No
20. (2019). Distinguished Lecture. Systems Research - Construed Broadly. Joint Duke University, University of North Carolina/Chapel Hill, NC State, North Carolina, United States of America
Invited?: Yes, Keynote?: No
21. (2019). Keynote. More than Storage. Mass Storage Systems and Technology(MSST), Santa Clara, United States of America
Invited?: Yes, Keynote?: Yes
22. Marie-Francoise, Roy Jessica Carter Anna Vasilchenko Anna Wienhard Fernando Seabra Chirigati. (2019). Implicit Bias Contributions to the Gender Gap in Science. Heidelberg Laureate Forum, Heidelberg, United States of America
Main Audience: Researcher
Invited?: Yes, Keynote?: No
23. (2018). Keynote. An NVM Carol. International Conference on Data Engineering, Paris, France
Main Audience: Researcher
Invited?: Yes, Keynote?: Yes
24. (2018). Flipped on Teaching. Harvard Club of Rochester Annual Meeting, Rochester, United States of America
Main Audience: General Public
Invited?: Yes, Keynote?: Yes
25. (2018). Distinguished Lecture. Systems Research - Construed Broadly. University of Washington, Washington, United States of America
Invited?: Yes, Keynote?: No
26. (2018). Guest Lecture. An NVM Carol. University of Santa Cruz Data Management course, Santa Cruz, United States of America
Invited?: Yes, Keynote?: No
27. (2018). Automatically Scalable Computation. DE Shaw Seminar, New York, United States of America
Main Audience: Researcher
Invited?: Yes, Keynote?: No
28. (2018). An NVM Carol. UBC Computer Science 50th Anniversary Celebration, Vancouver, Canada
Main Audience: General Public
Invited?: Yes, Keynote?: No
29. (2018). Distinguished Lecture. Automatically Scalable Computation. Johns Hopkins University, Maryland, United States of America
Invited?: Yes, Keynote?: No

30. (2018). Keynote. Automatically Scalable Computing. Israeli Systems Research Conference (SYSTOR), Haifa, Israel
Main Audience: Researcher
Invited?: Yes, Keynote?: Yes
31. (2017). Berkeley DB: The Good, The Bad, and the Ugly. Workshop on Failed Aspirations in Database Systems (FADS), Munich, Germany
Main Audience: Researcher
Invited?: Yes, Keynote?: No
32. (2017). Keynote. Automatically Scalable Computation. CodeMESH, London, United Kingdom
Main Audience: Researcher
Invited?: Yes, Keynote?: Yes
33. (2017). Distinguished Lecture Series. Automatically Scalable Computation. Northwestern Computer Science, Evanston, United States of America
Main Audience: Researcher
Invited?: Yes, Keynote?: No
34. (2017). Data Provenance: From Theory to Practice. Monthly Meeting of the IEEE Computer Society, Cambridge, United States of America
Main Audience: Knowledge User
Invited?: Yes, Keynote?: No
35. (2016). Automatically Scalable Computation. 2Sigma Distinguished Lecture Series, New York, United States of America
Main Audience: Researcher
Invited?: Yes, Keynote?: No
36. (2016). Careers in Academia. Women Engineers Code (WeCode), Cambridge, United States of America
Main Audience: General Public
Invited?: Yes, Keynote?: No
37. Mark Miller, David Mazières, Yuanyuan Zhou. (2015). Is achieving security a hopeless quest?. SOSP History Day, Monterey, United States of America
Main Audience: Researcher
Invited?: Yes, Keynote?: No
38. (2015). Automatically Scalable Computation. International Conference on Super Computing, Newport Beach, United States of America
Main Audience: Researcher
Invited?: Yes, Keynote?: Yes
39. Jonathan Zittrain, Yochai Benkler, Joseph Nai, Sophia Roosth. (2015). The New Cyber Infrastructure. Annual Meeting of the World Economic Forum, Davos, Switzerland
Main Audience: General Public
Invited?: Yes, Keynote?: No
40. (2013). World Domination Through Provenance. Workshop on Theory and Practice of Provenance, San Jose, United States of America
Main Audience: Researcher
Invited?: Yes, Keynote?: Yes
41. (2013). Keynote. Automatically Scalable Computation. RICON, New York, United States of America
Main Audience: Knowledge User
Invited?: Yes, Keynote?: Yes

42. Regina Herzlinger. (2013). Distinguished Lecture Series. Enabling Health Care Innovation through Technology: The Role of Academia. Washington University School Public Health, St. Louis, United States of America
Main Audience: General Public
Invited?: Yes, Keynote?: No
43. (2012). Provenance in Digital Collections. Meeting of the Library of Congress, Washington DC, United States of America
Main Audience: General Public
Invited?: Yes, Keynote?: No
44. Janice Ellig, Kristine Lilly, Patricia Florissi. (2012). Transform and Stay True. EMC World, Las Vegas, United States of America
Main Audience: General Public
Invited?: Yes, Keynote?: Yes

Publications

Journal Articles

1. Lerner, B., Boose, E., Ellison, A., Fong, E., Lau, M., Ngo, K., Pasquier, T., Perez, L., Seltzer, M., Sheehan, R., *Wonsil, J. (2023). Making Provenance Work for You. The R Journal.
Published,
Refereed?: Yes
2. *Wonsil, J., *Boufford, N., *Agrawal, P., *Chen, C., *Cui, T., *Sivaram, A., Seltzer, M. (2023). Reproducibility as a Service. Software: Practice and Experience.
Published,
Refereed?: Yes
3. *Wang, B., Seltzer, M. (2022). Tinkertoy: Build your own operating systems for IoT devices. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD), 10.1109/TCAD.2022.3198907, Best Paper Award.
Published,
Refereed?: Yes
4. *Hu, J., *Lu, E., *Holland, D., *Kawaguchi, M., Chong, S., Seltzer, M. (2022). Towards Porting Operating Systems with Program Synthesis. ACM Transactions on Programming Languages and Systems (TOPLAS).
Published,
Refereed?: Yes
5. Aliabadi, M., Seltzer, M., Asl, M., Ghavamizadeh, R. (2021). ARTINALI#: An Efficient Intrusion Detection Technique for Resource-Constrained Cyber-Physical Systems. International Journal of Critical Infrastructure. 100430(33)
Published, Elsevier,
Refereed?: Yes
6. Ellison, A., Boose, E., Lerner, B., Fong, E., Seltzer, M. (2020). People of Data: The End-to-End Provenance Project. IScience publication Patterns.
Published,
Refereed?: Yes
7. Gaynor, M., Schneider, D., Seltzer, M., Crannage, E., Barron, M., Waterman, J., Oberle, A. (2020). **A user-centered, learning asthma smartphone application for patients and providers.** Learning Health Systems.
Published,
Refereed?: Yes

8. Lau, M., Pasquier, T., Seltzer, M. (2020). **Rclean: A Tool for Writing Cleaner, More Transparent Code**. JOSS: The Journal of Open Source Software.
Published,
Refereed?: Yes
9. Mason, W. A., Doudali, T., Seltzer, M., Gavriloska, A. (2020). Unexpected Performance of Intel Optane™ DC Persistent Memory. IEEE Computer Architecture Letters.
Published,
Refereed?: Yes
10. Pasquier, T., Singh, J., Powles, J., Evers, D., Seltzer, M., Bacon, J. (2018). Data Provenance to Audit Compliance with Privacy Policy in the Internet of Things. Journal of Personal and Ubiquitous Computing. 22(2): 333-344.
Published,
Refereed?: Yes
11. Rao, M.*, Bacon, D.F., Parkes, D, Seltzer, M. (2018). Incentivizing Deep Fixes in Software Economics. IEEE Transactions on Software Engineering. : 21 pages.
Published,
Refereed?: Yes
12. Pasquier, T., Lau, M, Han*, X, Fong, E., Lerner, B., Boose, E, Crosas, M., Ellison, A., Seltzer, M. (2018). Sharing and Preserving Computations Analysis for Posterity with Encapsulator. IEEE Computing in Science and Engineering. 20(4): 111.
Published,
Refereed?: Yes
13. Angelino, E., Larus-Stone*, N, Alabi, D.*, Seltzer, M., Rudin, C. (2018). Learning Certifiably Optimal Rule Lists for Categorical Data. Journal of Machine Learning Research. 18(234): 1-78.
Published,
Refereed?: Yes
14. Pasquier, T., Lau, M., Trisovic, A., Boose, E., Couturier, B., Crosas, M., Ellison, A., Gibson, V., Jones, C., Seltzer, M. (2017). If these data could talk. Nature Scientific Data. 18: 5 pages.
Published,
Refereed?: Yes
15. Daniel Margo*, Margo Seltzer. (2015). A Scalable Distributed Graph Partitioner. Proceedings of the VLDB Endowment. 8(12): 1478-1489.
Published,
Refereed?: Yes
16. Carata, L., Akoush, S. Balakrishnan, N., Bytheway, T., Sohan, R., Seltzer, M., Hopper, A. (2014). A Primer on Provenance. Communications of the ACM 57. 5: 52-60.
Published,
Refereed?: Yes
17. Herzlinger, R., Seltzer, M., Gaynor, M. (2013). Applying KISS to Healthcare Information Technology. Computer. 46(11): 72-74.
Published,
Refereed?: Yes

Reports

1. David A. Holland*, Jinngei Hu*, Ming Kawaguchi, Eric Lu*, Stephen Chong, Margo Seltzer. (2019). Aquarium: Cassiopea and Alewife Languages. 20. arXiv.
2. Peter Kraft*, Amos Waterland*, Daniel Y Fu*, Anitha Gollamudi**, Shai Szulanski*, Margo Seltzer. (2018). Automatic Parallelization of Sequential Programs. 12. arXiv.

3. Hongyu Yang**, Cynthia Rudin, Margo Seltzer. (2017). Scalable Bayesian Rule Lists. 31. arXiv.

Conference Publications

1. (2023). Why write code when you can synthesize address translations?. Hot Topics in Operating Systems, Paper
Published
Refereed?: Yes, Invited?: Yes
2. (2023). Integrated Reproducibility with Self-describing Machine Learning Models. ACM Conference on Reproducibility, Paper
Published
Refereed?: Yes, Invited?: Yes
3. (2023). Optimal Sparse Regression Trees. Conference on Artificial Intelligence (AAAI), Paper
Published
Refereed?: Yes, Invited?: Yes
4. (2022). Tinkertoy: Build your own Operating System for IoT Devices. ACM SIGBED Conference on Embedded Systems (EMSOFT), Best Paper Award., Paper
Published
Refereed?: Yes, Invited?: Yes
5. (2022). Fast Sparse Decision Tree Optimization via Reference Ensembles. Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI-22), Paper
Published
Refereed?: Yes, Invited?: Yes
6. (2022). Arming IDS Researchers with a Robotic Arm Dataset. Dependable and Secure Networks Systems (DSN-2022), Paper
Published
Refereed?: Yes, Invited?: Yes
7. (2022). Anomaly Detection in Multiplex Dynamic Networks: from Blockchain Security to Brain Disease Prediction. NeurIPS workshop on Temporal Graph Learning, Select for workshop spotlight, Paper
Published
Refereed?: Yes, Invited?: Yes
8. (2022). TimberTrek: Exploring and Curating Trustworthy Decision Trees with Interactive Visualization. IEEE Visualization Conference (VIS-2022) (short paper), Paper
Published
Refereed?: Yes, Invited?: Yes
9. (2022). Fast Optimization of Weighted Sparse Decision Trees for use in Optimal Treatment Regimes and Optimal Policy Design. Advances in Interpretable Machine Learning and Artificial Intelligence Workshop (AIMLAI 2022), Paper
Published
Refereed?: Yes, Invited?: Yes

10. (2022). Classification for Generalized Linear and Additive Models. The 25th International Conference on Artificial Intelligence and Statistics (AISTATS-2022),
Paper
Published
Refereed?: Yes, Invited?: Yes
11. (2022). Shellac: Synthesis of a Multi-Pass Compiler. International Conference on Verified Software: Theories, Tools, and Experiments (VSTTE 2022),
Paper
Published
Refereed?: Yes, Invited?: Yes
12. (2022). Exploring the Whole Rashomon Set of Sparse Decision Trees. Thirty-Sixth Conference on Neural Information Processing (NeurIPS 2022), New Orleans LA USA, Selected for Oral Presentation,
Paper
Published
Refereed?: Yes, Invited?: Yes
13. (2022). FasterRisk: Fast and Accurate Interpretable Risk Scores. Thirty-Sixth Conference on Neural Information Processing (NeurIPS 2022),
Paper
Published
Refereed?: Yes, Invited?: Yes
14. (2021). ASSUAGE: Assembly Synthesis Using A Guided Exploration. User Interface Software Technology (UIST 2021),
Paper
Published
Refereed?: Yes, Invited?: Yes
15. (2020). Improving Data Scientist Efficiency with Provenance. International Conference on Software Engineering (ICSE),
Paper
Published
Refereed?: Yes, Invited?: Yes
16. (2020). Improving Data Scientist Efficiency with Provenance. Proceedings of the International Conference on Software Engineering. International Conference on Software Engineering, Seoul, Korea, North
Paper
Published
Refereed?: Yes, Invited?: No
17. (2020). UNICORN: Runtime Provenance-Based Detection for Advanced Persistent Threats. Network and Distributed System Security Symposium (NDSS),
Paper
Published
Refereed?: Yes, Invited?: Yes
18. (2020). **Towards Porting Operating Systems with Program Synthesis.** OOPSLA 2020,
Paper
Published
Refereed?: Yes, Invited?: Yes
19. (2020). Parking Packet Payload with P4. International Conference on emerging Networking and Experiments and Technologies (CoNEXT),
Paper
Published
Refereed?: Yes, Invited?: Yes

20. (2020). Smooth Kronecker: Solving the Combing Problem in Kronecker Graphs. Workshop on Graph Data Management Experiences (GRADES) and Network Data Analytics (NDA),
Paper
Published
Refereed?: Yes, Invited?: Yes
21. (2020). XANTHUS: Push-button Orchestration of Host Provenance Data Collection. Workshop on Practical and Reproducible Evaluation of Computer Systems (PREC-2020),
Paper
Published
Refereed?: Yes, Invited?: Yes
22. (2020). Babar: Revisiting Host-Based Intrusion Detection in the Age of Data Provenance. Network and Distributed System Security Symposium (NDSS),
Paper
Published
Refereed?: Yes, Invited?: No
23. (2020). Generalized and Scalable Optimal Sparse Decision Trees. Proceedings of the International Conference on Machine Learning (ICML-2020),
Paper
Published
Refereed?: Yes, Invited?: Yes
24. (2020). People of Data: The End-to-End Provenance Project. Patterns, DOI: <https://doi.org/10.1016/j.patter.2020.100016>,
Paper
Published
Refereed?: Yes, Invited?: Yes
25. (2020). SIGL: Securing Software Installations Through Deep Graph Learning. USENIX Security Symposium,
Paper
Published
Refereed?: Yes, Invited?: Yes
26. (2020). SIGL: Securing Software Installations Through Deep Graph Learning. Network and Distributed Systems Security Symposium, Seoul, Korea, North
Paper
Published
Refereed?: Yes, Invited?: No
27. (2019). ProvMark: A provenance expressiveness benchmarking system. Proceedings of the 20th ACM/IFIP International Middleware Conference (Middleware'19). 20th ACM/IFIP International Middleware Conference (Middleware'19),
Paper
Published
Refereed?: Yes, Invited?: No
28. (2019). Visionpaper – From Here to Provtopia. Proceedings of the 2019 Poly Workshop (Towards Polystores that manage multiple Databases, Privacy, Security, and/or Policy Issues for Heterogeneous Data),
Paper
Published
Refereed?: Yes, Invited?: No

29. (2019). Trials and Tribulations in Synthesizing Operating Systems. Proceedings of the Workshop on Programming Languages and Operating Systems. Workshop on Programming Languages and Operating Systems,
Paper
Published
Refereed?: Yes, Invited?: No
30. (2019). Optimal Sparse Decision Trees (spotlight paper; top 2.5%). Proceedings of the Conference on Neural Information Processing Systems (NeurIPS). Conference on Neural Information Processing Systems (NeurIPS), Vancouver,
Conference Date: 2019/12
Paper
Published
Refereed?: Yes, Invited?: No
31. (2018). Runtime Analysis of Whole-System Provenance. Proceedings of the 2018 Conference on Computer and Communications Security (CCS'18). Conference on Computer and Communications Security (CCS'18), Toronto, Canada (1601-1616)
Conference Date: 2018/10
Paper
Published
Refereed?: Yes, Invited?: Yes
32. (2018). Provenance-based Intrusion Detection: Opportunities and Challenges. Proceedings of the Workshop on the Theory and Practice of Provenance (TAPP 2018). Workshop on the Theory and Practice of Provenance (TAPP 2018), London, United Kingdom (4 pages)
Conference Date: 2018/7
Paper
Published
Refereed?: Yes, Invited?: Yes
33. (2018). Closing the Performance Gap Between Volatile and Persistent Key-Value Stores Using Cross-Referencing Logs. Proceedings of the 2018 USENIX Annual Technical Conference. 2018 USENIX Annual Technical Conference, Boston, United States of America (967-979)
Conference Date: 2018/6
Paper
Published
Refereed?: Yes, Invited?: No
34. (2017). Learning Certifiably Optimal Rule Lists for Categorical Data. Proceedings of the 23rd ACM Conference on Knowledge Discovery and Data Mining (KDD2017). 23rd ACM Conference on Knowledge Discovery and Data Mining (KDD2017), Halifax, Canada (35-44)
Conference Date: 2017/8
Paper
Published
Refereed?: Yes, Invited?: No
35. (2017). Scalable Bayesian Rule Lists. Proceedings of the International Conference on Machine Learning - Volume 70. International Conference on Machine Learning (ICML 2017), Sydney, Australia (3921-3930)
Conference Date: 2017/8
Paper
Published
Refereed?: Yes, Invited?: No

36. (2017). FRAPPuccino: Fault-detection through Runtime Analysis of Provenance. Workshop on Hot Topics in Cloud Computing (HotCloud 17), Santa Clara, United States of America
Conference Date: 2017/7
Paper
Published
Refereed?: Yes, Invited?: No
37. (2017). Practical Whole-System Provenance Capture. Proceedings of the 2017 Symposium on Cloud Computing. Symposium on Cloud Computing, Santa Clara, United States of America (405-418)
Conference Date: 2017/7
Paper
Published
Refereed?: Yes, Invited?: No
38. (2017). Persistent Memcached: Bringing Legacy Code to Byte-Addressable Persistent Memory. Workshop on Hot Topics in Storage and File Systems (HotStorage 17), Santa Clara, United States of America
Conference Date: 2015/7
Paper
Published
Refereed?: Yes, Invited?: No
39. (2017). A Crowdsourcing Approach to Collecting Tutorial Videos -- Toward Personalized Learning-at-Scale. Proceedings of the Fourth (2017) ACM Conference on Learning@ Scale, Cambridge, United States of America (157-160)
Conference Date: 2015/4
Paper
Published
Refereed?: Yes, Invited?: No
40. (2015). Towards General-Purpose Neural Network Computing. International Conference on Parallel Architectures and Compilation Techniques (PACT), Petrozavodsk, Russian Federation (99-112)
Conference Date: 2015/9
Paper
Published
Refereed?: Yes, Invited?: No
41. (2015). Recent Advances in Computer Architecture: The Opportunities and Challenges for Provenance. Workshop on the Theory and Practice of Provenance (TaPP), Edinburgh, United Kingdom
Conference Date: 2015/7
Paper
Published
Refereed?: Yes, Invited?: No
42. (2015). LLAMA: Efficient Graph Analytics Using Large Multiversed Arrays. 31st IEEE International Conference on Data Engineering (ICDE 2015), Seoul, Korea, South (363-374)
Conference Date: 2015/4
Paper
Published
Refereed?: Yes, Invited?: No
43. (2014). Programmable Smart Machines: A Hybrid Neuromorphic Approach to General Purpose Computation. Proceedings of Neuromorphic Architectures (NeuroArch) Workshop at 41th International Symposium on Computer Architecture (ISCA-41),
Paper
Published
Refereed?: Yes, Invited?: No

44. (2014). A Framework for Incentivizing Deep Fixes. WIT-EC. Workshop in Incentives and Trust in E-Commerce, Quebec City, Canada
Conference Date: 2014/7
Paper
Published
Refereed?: Yes, Invited?: No
45. (2014). Accelerating MCMC via parallel predictive prefetching. Conference on Uncertainty in Artificial Intelligence, Quebec City, Canada
Conference Date: 2014/7
Paper
Published
Refereed?: Yes, Invited?: No
46. (2014). ASC: Automatically Scalable Computation. ASPLOS. Conference on Architecture Support for Programming Languages and Operating Systems, Salt Lake City, United States of America
Conference Date: 2014/3
Paper
Published
Refereed?: Yes, Invited?: No
47. (2013). Evaluation of Filesystem Provenance Visualization Tools. Conference on Information Visualization, Atlanta, United States of America
Conference Date: 2013/10
Paper
Published
Refereed?: Yes, Invited?: No
48. (2013). Local Clustering in Provenance Graphs. ACM international conference on Information & Knowledge Management, Burlingame, United States of America (835-840)
Conference Date: 2013/10
Paper
Published
Refereed?: Yes, Invited?: No
49. (2013). Performance Introspection of Graph Databases. Proceedings of the 6th International Systems and Storage Conference, Haifa, Israel (18)
Conference Date: 2013/7
Paper
Published
Refereed?: Yes, Invited?: No
50. (2013). Computational Caches. Proceedings of the 6th International Systems and Storage Conference, Haifa, Israel (8)
Conference Date: 2013/7
Paper
Published
Refereed?: Yes, Invited?: No
51. (2013). Flash Caching on the Storage Client. USENIX ATC. USENIX Annual Technical Conference, San Jose, United States of America (127-138)
Conference Date: 2013/6
Paper
Published
Refereed?: Yes, Invited?: No

Intellectual Property

Patents

1. Committing copy-on-write transaction with a persist barrier for a persistent object including payload references. United States of America. 10229012.
Patent Status: Granted/Issued
Year Issued: 2019
2. Efficient memory management for persistent memory. United States of America. 15675528.
Patent Status: Granted/Issued
Year Issued: 2018
3. Efficient Copy-on-Write Transactions on Persistent Memory. United States of America. 15675533.
Patent Status: Granted/Issued
Year Issued: 2018
4. Data Structure Store in Persistent Memory. United States of America. 15489544.
Patent Status: Granted/Issued
Year Issued: 2018
5. Persistent memory transactions with undo logging. United States of America. 15675526.
Patent Status: Granted/Issued
Year Issued: 2018
6. Graph Processing using a Mutable Multilevel Graph Representation. United States of America. 9734607.
Patent Status: Granted/Issued
Year Issued: 2017