

RESEARCH INTERESTS

The design of data management systems, and the application of database concepts, to greatly extend the ability for domain-experts and normal users to work with data. Draws on database optimization, human-data-interaction, data provenance, visualization, and interface design.

Prepared April 6, 2023

EDUCATION

- Winter 2014 **Massachusetts Institute of Technology**, Cambridge, MA
Ph.D., Electrical Engineering and Computer Science
Advisor: Samuel Madden
Dissertation: Explaining Data in Visual Analytic Systems
- May 2010 **Massachusetts Institute of Technology**, Cambridge, MA
M.S., Electrical Engineering and Computer Science
Advisor: Samuel Madden
Dissertation: Shinobi: Insert-aware Partitioning and Indexing Techniques For Skewed Database Workloads
- Spring 2007 **UC Berkeley**, Berkeley, CA
B.S., Electrical Engineering and Computer Science

PROFESSIONAL EXPERIENCE

- 2021 Spring On parental leave
- 2020–current **Columbia University**, NY, NY
Associate Professor (untenured) – Computer Science
Co-director – Center for Data, Media & Society, Columbia Data Science Institute
Co-advisor – Columbia Journalism & Computer Science Dual Degree Program
- 2015–2020 **Columbia University**, NY, NY
Assistant Professor – Computer Science
Co-director – Center for Data, Media & Society, Columbia Data Science Institute
Co-advisor – Columbia Journalism & Computer Science Dual Degree Program
- 2015 **UC Berkeley**, Berkeley, CA
Visiting Scholar – AMPLab
- 2008–2014 **Massachusetts Institute of Technology**, Cambridge, MA
Ph.D. Student – CSAIL
- 2007–2008 **Google Inc.**, Mountain View, CA
Research Intern – Data Management Group

AWARDS

- 2023 **Adobe Data Science Research Award**
- 2021 **Google Research Scholar Award**
Amazon Faculty Award
- 2020 **Very Large Data Bases (VLDB) Distinguished Reviewers Award**
- 2019 **Conference on Innovative Data Systems Research (CIDR) Best Gong-show Award** for best 5-minute short talk award.
SIG Management of Data (SIGMOD) Exceptional Service Award
NSF CAREER Award
- 2018 **Very Large Data Bases (VLDB) 10 Year Test-of-Time Award** for the WebTable project, which analyzed over 14 billion HTML tables to collect the largest corpus of database-like tables on the web.
Google Faculty Award for development of deep neural inspection, to identify high-level logic in neural network models.
Amazon Faculty Award for development of human-in-the-loop techniques to understand machine learning models.
- 2016 **SIG Management of Data (SIGMOD) best demo award** for demonstration of model-aware database cleaning techniques.

GRANTS

Total Awarded \$3,095,036

Awarded

- 2023 Adobe: Interactive Explanatory Analyses for Big Data Business KPIs
PI: Eugene Wu.
- 2021 Columbia SIRS: Filling “Data Voids” to Protect the Food Supply Against Climate Change (Renewal)
PI: Eugene Wu. Co-PI: Lydia Chilton. Co-PI: Daniel Osgood, Columbia University Earth Institute
- NSF Elements: Decision Engine for Socioeconomic Disaster Risk (DESDR) - Data Collection, Fusion, and Analysis to Protect Vulnerable Populations From Extreme Weather
PI: Daniel Osgood. Co-PI: Lydia Chilton. Co-PI: Eugene Wu
- Google Research Scholar Award
PI: Eugene Wu.
- Amazon Faculty Award
PI: Eugene Wu.
- NSF III: Medium: Unified Prefetch Framework for Approximation Tolerant Interactive Applications
PI: Eugene Wu. Co-PI: Dan Rubenstein, Columbia University; Co-PI: Ravi Netravali, UCLA
- 2020 NSF REU: CAREER: Visual Database Interfaces
PI: Eugene Wu
- NSF III: Small: Bringing database query optimization to data intensive applications
PI: Ken Ross. Co-PI: Eugene Wu
- Columbia SIRS: Filling “Data Voids” to Protect the Food Supply Against Climate Change
PI: Eugene Wu. Co-PI: Lydia Chilton. Co-PI: Daniel Osgood, Columbia University Earth Institute
- 2019 NSF CAREER: Visual Database Interfaces
PI: Eugene Wu
- 2018 Google Faculty Award
PI: Eugene Wu
- Amazon Faculty Award
PI: Eugene Wu
- Columbia DSI Scholars: Data Cleaning, Data Reconciliation and Data Visualization
PI: Eugene Wu. Co-PI: Daniel Osgood, Columbia University Earth Institute
- 2017 Columbia Alliance: Perceptual Functions for Faster Interactive Visualizations
PI: Eugene Wu. Co-PI: Yanlei Diao, Ecole-Polytechnique

NSF: I-Corps: Internet of Things Monitoring System
PI: Eugene Wu.

2016 NSF: ACM SIGMOD Conference 2016: Student Activities and Travel Support
PI: Eugene Wu.

NSF IIS: Medium: Collaborative Research: Composing Interactive Data Visualizations
PI: Joseph Hellerstein, Berkeley. Co-PI: Jeff Heer, University of Washington. Co-PI: Eugene Wu.

2015 NSF III: Small: Collaborative Research: Towards Interactive Data Visualization Management Systems
PI: Arnab Nandi, Ohio State University. Co-PI: Eugene Wu.

NSF REU: Development of Graphical Perception as a Service
PI: Eugene Wu.

Glossary

NSF National Science Foundation
RI Robust Intelligence
III Information Integration and Informatics
HCC Human-Centered Computing
REU Research Experiences for Undergraduates

PUBLICATIONS

H-index: 34.

Total citations: 6,782.

Citation counts are reported from <http://scholar.google.com> when exceeding ten citations. Acceptance rates are reported for refereed conference papers when available.

Authorship Conventions

Conferences are the primary publication venue in data management. Senior authors are typically listed last. Advisees are listed in **bold**. Eugene Wu is underlined.

In Progress or In Review

- 2023 [1] **L. Ramjit**, **Y. Chen**, S. Mitra, R. Netravali, E. Wu. “Physical Visualization Design.” In Progress 2023.
- [2] **H. Mohammed**, **C. Summers**, E. Wu. “SmokedDuck: Fast Lineage in Analytical Databases.” In Progress 2023.
- [3] **Z. Huang**, R. Sen, **J. Liu**, E. Wu. “JoinBoost: Grow Trees Over Normalized Data Using Only SQL.” In Review 2023.
- [4] **Z. Huang**, **J. Liu**, **D. Alabi**, R. C. Fernandez, E. Wu. “Saibot: A Differentially Private Data Search Platform.” In Review 2023.
- [5] **Z. Huang**, **P. Kalyan**, R. Sen, E. Wu. “Random Forests over Normalized Data in CPU-GPU DBMSes.” In Review 2023.
- [6] **Z. Huang**, **P. Kalyan**, E. Wu. “Analysis Errors Over Semantic Layers and How To Avoid Them.” In Progress 2023.
- [7] **J. Liu**, **Z. Huang**, E. Wu. “Demonstration of Kitana: Interactive and Private Data Augmentation Search for AutoML.” In Progress 2023.
- [8] **Y. Jiang**, C. Xiong, S. Franconeri, E. Wu. “Comparison Visualization Under Data-induced Channels.” In Progress 2023.
- [9] **Y. Chen**, **J. Tao**, E. Wu. “Data Interface Grammar: An Abstraction for Interactive Visualization Interfaces.” In Progress 2023.
- [10] A. Suh, A. Mosca, E. Wu, R. Chang. “A Hypothesis-Based Framework for Evaluating Visualization and Visual Analytics Systems.” In Review 2023.
- [11] **A. Yao**, **L. Flokas**, E. Wu. “FaDE: Answering “Why?” Made Fast.” In Progress 2023.
- [12] **Z. Huang**, P. Subramaniam, R. Fernandez, E. Wu. “Kitana: A Data-as-a-Service Platform.” In Review 2023.
- [13] **Z. Huang**, E. Wu. “Calibration: A Simple Trick for Fast Interactive Join Analytics.” In Review 2023.

Full Publications

- [14] Y. Wang, Y. Wang, X. Chen, Y. Zhao, F. Zhang, E. Wu, C. Fu, X. Yu. “OM3: An Ordered Multi-level Min-Max Representation for Interactive Progressive Visualization of Time Series.” SIGMOD 2023.
- [15] G. Vitagliano, M. Hameed, L. Jiang, L. Reisener, E. Wu, F. Naumann. “Pollock: A Data Loading Benchmark.” VLDB 2023.
- [16] X. Wang, W. Wu, J. Wu, Y. Chen, N. Zrymiak, C. Qu, **L. Flokas**, G. Chow, J. Wang, T. Wang, E. Wu, Q. Zhou. “ConnectorX: Accelerating Data Loading From Databases to Dataframes.” VLDB 2023.
- 2022 [17] A. Suh, **Y. Jiang**, A. Mosca, E. Wu, R. Chang. “A Grammar for Hypothesis-Driven Visual Analysis.” ArXiv 2022.
- [18] **L. Flokas**, W. Wu, J. Wang, N. Verma, E. Wu. “How I Stopped Worrying About Training Data Bugs and Started Complaining.” DEEM Workshop 2022.
- [19] **J. Tao**, **Y. Chen**, E. Wu. “Interactive Interface Generation in Notebooks.” SIGMOD 2022 demo.
- [20] **Y. Chen**, E. Wu. “PI2: Generating Visual Analysis Interfaces From Queries.” SIGMOD 2022.
- [21] E. Wu. “View Composition Algebra for Ad Hoc Comparisons.” TVCG 2022.
- [22] **Z. Huang**, E. Wu. “Reptile: Aggregation-level Explanations for Hierarchical Data.” SIGMOD 2022.
- [23] I. Drori, S. Tran, R. Wang, N. Cheng, K. Liu, L. Tang, E. Ke, N. Singh, T. L. Patti, J. Lynch, A. Shporer, N. Verma, E. Wu, G. Strang. “A Neural Network Solves and Generates Mathematics Problems by Program Synthesis: Calculus, Differential Equations, Linear Algebra, and More.” PNAS 2022 (in review).
- [24] **L. Flokas**, Y. Wu, J. Wang, N. Verma, E. Wu. “Complaint-Driven Training Data Debugging at Interactive Speeds.” SIGMOD 2022.
- [25] Y. Wu, Y. Liu, **L. Flokas**, J. Wang, E. Wu. “Enabling SQL-based training data debugging for federated learning.” VLDB 2022.
- 2021 [26] B. Lockhard, J. Wang, E. Wu. “Explaining SQL-ML Queries with Bayesian Optimization.” VLDB 2021.
- [27] Y. Wu, R. Chang, J. Hellerstein, A. Satyanarayan, E. Wu. “DIEL: Interactive Visualization Beyond the Here and Now.” VIS 2021.
- [28] **R. Netzorg**, **L. Arnett**, A. Chaintreau, E. Wu. “PopFactor: Live-Streamer Behavior and Popularity.” ICWSM 2021. (acceptance rate: approx 20%)
- [29] M. Procopio, A. Mosca, C. Scheidegger, E. Wu, R. Chang. “Impact of Cognitive Biases on Progressive Visualization.” TVCG 2021. (acceptance rate: 24.6%)
- [30] F. Neutatz, B. Chen, Z. Abedjan, E. Wu. “From Cleaning Before ML to Cleaning For ML.” Invited, IEEE Data Engineering Bulletin 2021.
- 2020 [31] **H. Mohammed**, Z. Wei, R. Netravali, E. Wu. “Continuous Prefetch for Interactive Data Applications.” VLDB 2020. (acceptance rate: 24.8%)
- [32] Y. Wu, **L. Flokas**, J. Wang, E. Wu. “Complaint-driven Training Data Debugging for Query 2.0.” SIGMOD 2020. (acceptance rate: 26.9%)
- [33] **Y. Chen**, E. Wu. “Monte Carlo Tree Search for Generating Interactive Data Analysis Interfaces.” Intelligent Process Automation (IPA) 2020.

- 2019 [34] **L. Ramjit**, M. Interlandi, **E. Wu**, R. Netravali. “*Acorn: Aggressive Result Caching in Spark SQL.*” SOCC 2019. (acceptance rate: 25%)
- [35] N. Tang, **E. Wu**, G. Li. “*Towards Democratizing Relational Data Visualization.*” SIGMOD 2019 Tutorial.
- [36] **Q. Zhang**, **H. Zhang**, **V. Rai**, **T. Sellam**, **E. Wu**. “*Precision Interfaces.*” SIGMOD 2019. (acceptance rate: 20%)
- [37] P. Wang, J. Wang, R. Shea, **E. Wu**. “*Progressive Deep Web Crawling Through Keyword Queries For Data Enrichment.*” SIGMOD 2019. (acceptance rate: 20%)
- [38] **T. Sellam**, **K. Lin**, **I. Y. Huang**, M. Yang, C. Vondrick, **E. Wu**. “*DeepBase: Deep Inspection of Neural Networks.*” SIGMOD 2019. (acceptance rate: 20%, citations: 11)
- 2018 [39] M. Cafarella, A. Halevy, D. Z. Wang, H. Lee, J. Madhavan, C. Yu, **E. Wu**. “*Ten Years of Web Tables.*” PVLDB 2018 Invited Paper. (citations: 32)
- [40] **G. Ryan**, A. Mosca, R. Chang, **E. Wu**. “*At a Glance: Approximate Entropy as a Measure of Line Chart Visualization Complexity.*” InfoVIS 2018. (acceptance rate: 25.13%, citations: 13)
- [41] **H. Nilforoshan**, **E. Wu**. “*Leveraging Quality Prediction Models for Automatic Writing Feedback.*” ICWSM 2018. (acceptance rate: 20%)
- [42] **T. Sellam**, **K. Lin**, **I. Y. Huang**, C. Vondrick, **E. Wu**. “*“I Like the Way You Think!” Inspecting the Internal Logic of Recurrent Neural Networks.*” SysML 2018.
- [43] **F. Psallidas**, **E. Wu**. “*Smoke: Fine-grained Lineage at Interactive Speeds.*” VLDB 2018. (acceptance rate: 21%, citations: 43)
- 2017 [44] L. Sun, M. J. Franklin, J. Wang, **E. Wu**. “*Skipping-oriented Partitioning for Columnar Layouts.*” VLDB 2017. (acceptance rate: 18.6%, citations: 26)
- [45] **E. Wu**, **F. Psallidas**, **Z. Miao**, **H. Zhang**, L. Rettig, Y. Wu, **T. Sellam**. “*Combining Design and Performance in a Data Visualization Management System.*” CIDR 2017. (citations: 26)
- [46] X. Wang, A. Meliou, **E. Wu**. “*QFix: Diagnosing errors through query histories.*” SIGMOD 2017. (citations: 25)
- 2016 [47] D. Haas, J. Wang, **E. Wu**, M. J. Franklin. “*CLAMShell: Speeding up Crowds for Low-latency Data Labeling.*” VLDB 2016. (acceptance rate: 35.5%, citations: 65)
- 2015 [48] A. A. Bhattacharya, D. Hong, D. Culler, J. Ortiz, K. Whitehouse, **E. Wu**. “*Automated Metadata Construction to Support Portable Building Applications.*” BuildSys 2015. (citations: 64)
- [49] **E. Wu**. “*Explaining Data in Visual Analytic Systems.*” Doctoral Thesis 2015.
- 2014 [50] **E. Wu**, L. Battle, S. Madden. “*The Case for Data Visualization Management Systems.*” VLDB 2014. (citations: 85)
- [51] A. Jindal, P. Rawlani, **E. Wu**, S. Madden, A. Deshpande, M. Stonebraker. “*Vertexica: Your Relational Friend for Graph Analytics!*” SIGMOD 2014 demo. (citations: 63)
- 2013 [52] A. Cheung, L. Ravindranath, **E. Wu**, S. Madden, H. Balakrishnan. “*Mobile applications need Targeted Micro-updates.*” APSYS 2013.
- [53] **E. Wu**, S. Madden. “*Scorpion: Explaining Away Outliers in Aggregate Queries.*” VLDB 2013 (Best-of). (acceptance rate: 22.7%, citations: 211)
- [54] **E. Wu**, S. Madden, M. Stonebraker. “*SubZero: a Fine-Grained Lineage System for Scientific Databases.*” ICDE 2013 (Best-of). (citations: 43)
- 2012 [55] A. Marcus, **E. Wu**, D. Karger, S. Madden, R. Miller. “*Human-powered Sorts and Joins.*” VLDB 2012. (acceptance rate: 20.3%, citations: 334)

- 2011 [56] E. Wu, S. Madden. “*Partitioning Techniques for Fine-Grained Indexing.*” ICDE 2011. (acceptance rate: 20%, citations: 39)
- [57] A. Marcus, E. Wu, D. Karger, S. Madden, R. Miller. “*Demonstration of Qurk: A Query Processor for Human Operators.*” SIGMOD 2011. (citations: 51)
- [58] E. Wu, C. Curino, S. Madden. “*No Bits Left Behind.*” CIDR 2011.
- [59] A. Marcus, E. Wu, S. Madden, R. Miller. “*Crowdsourced Databases: Query Processing with People.*” CIDR 2011. (citations: 248)
- [60] C. Curino, E. Jones, R. Popa, N. Malviya, E. Wu, S. Madden, H. Balakrishnan, N. Zeldovich. “*Relational Cloud: A Database-as-a-Service for the Cloud.*” CIDR 2011. (citations: 484)
- 2010 [61] P. Cudre-Mauroux, E. Wu, S. Madden. “*TrajStore: An Adaptive Storage System for Very Large Trajectory Data Sets.*” ICDE 2010. (citations: 214)
- 2009 [62] P. Cudre-Mauroux, E. Wu, S. Madden. “*The Case for RodentStore: An Adaptive, Declarative Storage System.*” CIDR 2009. (citations: 32)
- 2008 [63] M. Cafarella, A. Halevy, D. Wang, E. Wu, Y. Zhang. “*WebTables: Exploring the Power of Tables on the Web.*” VLDB 2008. (acceptance rate: 16.8%, citations: 737)
- 2006 [64] E. Wu, Y. Diao, S. Rizvi. “*High-performance complex event processing over streams.*” SIGMOD 2006. (acceptance rate: 13%, citations: 1138)
- 2005 [65] M. J. Franklin, S. R. Jeffery, S. Krishnamurthy, F. Reiss, S. Rizvi, E. Wu, O. Cooper, A. Edakkunni, W. Hong. “*Design Considerations for High Fan-In Systems: The HiFi Approach.*” CIDR 2005. (citations: 268)

Short Papers, Technical Reports, and Demos

- 2022 [66] Y. Chen, R. Li, A. Mac, T. Xie, T. Yu, E. Wu. “*NL2INTERFACE: Interactive Visualization Interface Generation from Natural Language Queries.*” VIS nlviz workshop 2022.
- [67] S. Cheng, H. Zhu, E. Wu. “*How Do Captions Affect Visualization Reading?.*” VIS Viscomm 2022.
- [68] E. Wu. “*Extending the View Composition Algebra to Hierarchical Data.*” arXiv 2022.
- [69] Y. Liu, M. Zhao, S. Xia, E. Wu, X. Jiang. “*A Sensorless Drone-based System for Mapping Indoor 3D Airflow Gradients.*” MobiSys 2022 Demo.
- 2021 [70] J. Fisher, R. Chang, E. Wu. “*Dynamic Breakpoints for Y-axis Scales.*” InfoVIS 2021 (short paper).
- 2020 [71] Y. Wu, R. Chang, J. Hellerstein, E. Wu. “*Facilitating Exploration with Interaction Snapshots under High Latency.*” InfoVIS (short paper) 2020.
- [72] L. Zhao, Q. Li, P. Wang, J. Wang, E. Wu. “*ActiveDeeper: A Model-based Active Data Enrichment system.*” VLDB 2020 demo.
- [73] L. Ramjit, Z. Kong, R. Netravali, E. Wu. “*Physical Visualization Design.*” SIGMOD (demo) 2020.
- [74] L. Flokas, Y. Wu, J. Wang, E. Wu. “*Towards Complaint-driven ML Workflow Debugging.*” MLOps 2020.
- 2019 [75] S. Krishnan, E. Wu. “*AlphaClean: Automatic Generation of Data Cleaning Pipelines.*” ArXiv 2019. (citations: 12)
- [76] L. Arnett, R. Netzorg, A. Chaintreau, E. Wu. “*Cross-platform Interactions and Popularity in the Live-streaming Community.*” CHI Latebreaking 2019.

- [77] E. Wu. “CIDR2: Crazier Innovations in Databases JOIN Reinforcement-learning Research.” CIDR 2019 Abstract.
- 2018 [78] Y. Chen, Y. Shi, B. Chen, T. Sellam, C. Vondrick, E. Wu. “Deep Neural Inspection Using Deep-Base.” LearnSys 2018 Workshop at NIPS.
- [79] F. Psallidas, E. Wu. “Provenance in Interactive Visualizations.” HILDA 2018.
- [80] H. Zhang, V. Rai, T. Sellam, E. Wu. “Precision Interfaces for Different Modalities.” SIGMOD (demo) 2018.
- [81] F. Psallidas, E. Wu. “Demonstration of Smoke: A Deep Breath of Data-Intensive Lineage Applications.” SIGMOD (demo) 2018.
- [82] P. Wang, Y. He, R. Shea, J. Wang, E. Wu. “Deeper: A Data Enrichment System Powered by Deep Web.” SIGMOD (demo) 2018.
- 2017 [83] H. Zhang, T. Sellam, E. Wu. “Mining Precision Interfaces From Query Logs.” Tech Report 2017.
- [84] S. Krishnan, M. J. Franklin, K. Goldberg, E. Wu. “BoostClean: Automated Error Detection and Repair for Machine Learning.” Tech Report 2017. (citations: 36)
- [85] M. Procopio, C. Scheidegger, E. Wu, R. Chang. “Load-n-Go: Fast Approximate Join Visualizations That Improve Over Time.” DSIA 2017.
- [86] G. Ryan, A. Mosca, E. Wu, R. Chang. “Approximate Entropy as a Measure of Line Chart Complexity.” InfoVIS Poster 2017.
- [87] Y. Wu, L. Xu, R. Chang, E. Wu. “Towards a Bayesian Model of Data Visualization Cognition.” DECISIVE 2017.
- [88] H. Nilforoshan, J. Wang, E. Wu. “PreCog: Improving Crowdsourced Data Quality Before Acquisition.” Arxiv 2017.
- [89] H. Zhang, T. Sellam, E. Wu. “Precision Interfaces.” HILDA 2017.
- [90] S. Krishnan, E. Wu. “PALM: Machine Learning Explanations For Iterative Debugging.” HILDA 2017. (citations: 42)
- [91] H. Nilforoshan, J. Sands, K. Lin, R. Khanna, E. Wu. “Segment-Predict-Explain for Automatic Writing Feedback.” Collective Intelligence 2017.
- [92] H. Nilforoshan, J. Sands, K. Lin, R. Khanna, E. Wu. “Dialectic: Enhancing Text Input Fields with Automatic Feedback to Improve Social Content Writing Quality.” ArXiv 2017.
- [93] E. Wu. “CIDR: Chat-oriented Innovations in Database Research.” CIDR 2017 Abstract.
- 2016 [94] Y. Wu, J. Hellerstein, E. Wu. “A DeVIL-ish Approach to Inconsistency in Interactive Visualizations.” HILDA 2016.
- [95] D. Alabi, E. Wu. “PFunk-H: Approximate Query Processing using Perceptual Models.” HILDA 2016. (citations: 21)
- [96] S. Krishnan, D. Haas, M. J. Franklin, E. Wu. “Towards Reliable Interactive Data Cleaning: A User Survey and Recommendations.” HILDA 2016. (citations: 47)
- [97] N. Kamat, E. Wu, A. Nandi. “TrendQuery: A System for Interactive Exploration of Trends.” HILDA 2016.
- [98] S. Krishnan, M. Franklin, K. Goldberg, J. Wang, E. Wu. “ActiveClean: An Interactive Data Cleaning Framework For Modern Machine Learning.” SIGMOD 2016 Demo. (citations: 119)
- [99] E. Wu, L. Jiang, L. Xu, A. Nandi. “Graphical Perception in Animated Bar Charts.” Arxiv 2016. (citations: 12)

- [100] X. Wang, A. Meliou, E. Wu. “*QFix: Demonstrating error diagnosis in query histories.*” SIGMOD 2016 Demo.
- [101] X. Wang, A. Meliou, E. Wu. “*QFix: Diagnosing errors through query histories.*” Arxiv 2016.
- [102] S. Krishnan, J. Wang, E. Wu, M. J. Franklin, K. Goldberg. “*ActiveClean: Interactive Data Cleaning While Learning Convex Loss Models.*” Arxiv 2016.
- [103] L. Battle, E. Benson, A. Parameswaran, E. Wu. “*Indexing Cost Sensitive Prediction.*” Technical Report 2016.
- 2015 [104] E. Wu, A. Nandi. “*Towards Perception-aware Interactive Data Visualization Systems.*” DSIA 2015. (citations: 14)
- [105] S. Krishnan, J. Wang, M. J. Franklin, K. Goldberg, T. Kraska, T. Milo, E. Wu. “*SampleClean: Fast and Reliable Analytics on Dirty Data (overview paper).*” IEEE Data Eng. Bulletin 2015. (citations: 34)
- [106] D. Haas, S. Krishnan, J. Wang, M. J. Franklin, E. Wu. “*Wisteria: Nurturing Scalable Data Cleaning Infrastructure.*” VLDB 2015 demo. (citations: 39)
- [107] A. Bhardwaj, A. Deshpande, A. Elmore, D. Karger, S. Madden, A. Parameswaran, H. Subramanyam, E. Wu, R. Zhang. “*Collaborative Data Analytics with Datahub.*” VLDB 2015 demo. (citations: 54)
- 2013 [108] E. Wu, A. Marcus, S. Madden. “*Data In Context: Aiding News Consumers while Taming Datas-paces.*” DBCrowd 2013.
- 2012 [109] E. Wu, S. Madden, M. Stonebraker. “*A Demonstration of DBWipes: Clean as You Query.*” VLDB 2012.
- 2010 [110] C. Curino, E. Jones, Y. Zhang, E. Wu, S. Madden. “*Relational Cloud: The Case for a Database Service.*” MIT Tech Report 2010.
- 2009 [111] E. Wu, P. Cudre-Mauroux, S. Madden. “*Demonstration of the TrajStore System.*” VLDB 2009 demo.
- 2008 [112] M. Cafarella, N. Khoussainova, D. Wang, E. Wu, Y. Zhang, A. Halevy. “*Uncovering the Relational Web.*” WebDB 2008. (citations: 178)
- 2007 [113] D. Gyllstrom, E. Wu, H. Chae, Y. Diao, P. Stahlberg, G. Anderson. “*SASE: Complex Event Processing over Streams (Demo).*” CIDR 2007. (citations: 276)
- 2006 [114] D. Gyllstrom, E. Wu, H. Chae, Y. Diao, P. Stahlberg, G. Anderson. “*SASE: Complex Event Processing over Streams.*” CoRR 2006.
- [115] M. N. Garofalakis, K. P. Brown, M. J. Franklin, J. M. Hellerstein, D. Z. Wang, E. Michelakis, L. Tancau, E. Wu, S. R. Jeffery, R. Aipperspach. “*Probabilistic Data Management for Pervasive Computing: The Data Furnace Project.*” IEEE Data Eng. Bulletin 2006. (citations: 41)
- 2004 [116] O. Cooper, A. Edakkunni, M. J. Franklin, W. Hong, S. R. Jeffery, S. Krishnamurthy, F. Reiss, S. Rizvi, E. Wu. “*HiFi: A Unified Architecture for High Fan-in Systems.*” VLDB 2004 Demo. (citations: 39)

Primary Conference Names and 5 Year Impact Factor

VLDB	Very Large Databases. 3.56.
SIGMOD	SIG Management of Data. 3.41.
ICDE	IEEE International Conference on Data Engineering. 2.63.
CIDR	Conference on Innovative Data system Research. 3.3.
SOCC	ACM Symposium on Cloud Computing. 4.79.

TVCG Transactions on Visualization and Computer Graphics. 4.558.
InfoVIS Information Visualization. 4.558.
EuroVIS Eurographics/IEEE Symposium on Visualization. 2.15
HILDA Human in the Loop Data Analysis Workshop. 2.0
BuildSys Conference on Systems for Energy-Efficient Built Environments. 2.02
DSIA Data Systems for Interactive Analysis Workshop.
SysML Conference on Machine Learning and Systems.

TEACHING EXPERIENCE

At Columbia University

- 2023 Spring *Instructor, Topics in Database Research (W6113, approx 15 Students)*
<https://w6113.github.io/>
- 2022 Fall *Instructor, Introduction to Databases (W4111, 123 Students)*
<https://w4111.github.io/>
- 2022 Spring *Instructor, Introduction to Databases (W4111, 181 students)*
<https://w4111.github.io/>
- 2021 Fall *Instructor, Systems for Human Data Interaction (W6998, 15 students)*
<https://columbiaviz.github.io/>
- 2020 Fall *Instructor, Topics in Database Research (W6113, 15 students)*
<http://w6113.github.io>
- 2020 Spring *Instructor, Systems for Human Data interaction (W6998, 13 students)*
https://columbiaviz.github.io/2020s_w6998/
- 2019 Spring *Instructor, Introduction to Databases (W4111, 73 students)*
<http://w4111.github.io>
- 2019 Spring *Instructor, Database Research Topics (W6113, 13 students)*
<http://w6113.github.io>
- 2018 Fall *Instructor, Introduction to Databases (W4111, 107 students)*
<http://w4111.github.io>
- 2018 Spring *Instructor, Database Topics in Research & Practice (W6998.005, 9 students)*
<http://columbiadb.github.io>
- 2018 Spring *Co-Instructor, Computing Systems for Data Science (W4121, 158 students)*
<http://w4121.github.io>
- 2017 Spring *Instructor, Interactive Data Exploration Systems (W6998.002, 10 students)*
<http://columbiaviz.github.io>
- 2017 Spring *Co-Instructor, Computing Systems for Data Science (W4121, 124 students)*
<http://w4121.github.io>
- 2016 Fall *Instructor, Introduction to Databases (W4111, 139 students)*
<http://w4111.github.io>
- 2016 Fall *Instructor, Introduction to Databases CVN sec. (W4111, 39 students)*
<http://w4111.github.io>
- 2016 Spring *Co-Instructor, Big Data Systems (W4121, 79 students)*
<http://w4121.github.io>
- 2015 Fall *Instructor, Introduction to Databases (W4111, 108 students)*

<http://w4111.github.io>

Prior to Columbia University

- 2013 Fall *Instructor, From Ascii To Answers (MIT 6.885)*
 I co-developed and instructed MIT's first Big Data course focused on large scale data analysis tools and techniques. Topics ranged from data cleaning and integration, large-scale systems like Hadoop, to scalable visualization techniques. We developed eight labs to give students hands-on experience with the systems covered in class. The course is freely available online at <http://github.com/mitdbg/asciiclass>
- 2012 Spring *Instructor, Introduction to Data Analysis*
 I co-developed and taught an Introduction to Data Analysis course to approximately 20 students during MIT's Independent Activities Period in January. The course is freely available online at <http://dataiap.github.io>
- 2011 – 2012 *Head of Curriculum, MEET*
 MEET is a 3-year technology program and peace initiative that teaches Israeli and Palestinian high school students. I organized curriculum preparation for each year's incoming instructors. I also successfully migrated the organization from a Java-based curriculum to a Python-oriented one and developed the lesson plans for the transition.
- 2010 Fall *Teaching Assistant, Database Systems (MIT 6.830)*
 I assisted in writing and grading the assignments and projects.
- 2010 Summer *Instructor, MEET*
 I mentored a group of 30 Israeli and Palestinian high school students as part of the MIT MEET program, a peace initiative in the Middle East centered around teaching computer science.
- 2010, 2011 *Instructor, Introduction to Java Course (MIT 6.S092)*
 Spring I instructed a class of 50 students in an introduction to the Java programming language. MIT does not have such an introductory course, so this course is taken by many MIT undergraduates to prepare them for 6.004, a core course that assumes proficiency in Java. The course is freely available online at <http://bit.ly/alvK9m>
- 2006 Fall *Teaching Assistant, Database Systems (UCB CS186)*
 I taught approximately 30 students in weekly discussion sections. I assisted in writing and grading the assignments and projects.

STUDENTS

- Postdocs Thibault Sellam 2017-2018. Now at Google.
- Graduated PhDs Fotis Psallidas 2016-2019. Now at Microsoft Research Gray Systems Lab
Thesis: Physical Plan Instrumentation in Databases: Design Principles and Applications
- Current PhDs Yiru Chen 2018-current
Google Research Fellowship 2021
Lampros Flokas 2018-2022 expected
Haneen Mohammed 2018-current
PhD Winner, ACM SIGMOD Student Research Competition 2020
Zachary Huang 2019-current
Charles Summers 2021-current
Jerry Liu 2022-current
- PhD Committee at Columbia Orestis Polychroniou. Defended 2017
Thesis: Analytical Query Execution Optimized for all Layers of Modern Hardware
Ioannis Paparrizos. Defended 2018
Thesis: Fast, Scalable, and Accurate Algorithms for Time-Series Analysis
Daniel Miao. Defended 2018
Thesis: Personalized Navigational Instruments for Map User Interfaces
Wangda Zhang. Defended 2020
Thesis: Optimizing Query Processing Under Skew
- External Mohammad Mahdavi, TU Berlin. Defended 2020. Primary advisor: Ziawasch Abedjan
Marianne Procopio, Tufts. Defended 2020. Primary advisor: Remco Chang
Yifan Wu, UC Berkeley. Defended 2021. Primary advisor: Joe Hellerstein
Lana Ramjit, UCLA. Defended 2021. Primary advisor: Ravi Netravali
- MS Advising Daniel Alabi 2015-2016. Now Ph.D. at Harvard University
Zhengjie Miao 2016-2017. Now Ph.D. at Duke University
Gabriel Ryan 2017-2018. Now Ph.D. at Columbia University
Haoci Zhang 2016-2018. Now at Facebook
Tejas Dharamsi 2017-2018. Now at Trifacta
Sharan Suryanarayanan 2018. Now at Facebook
Charles Summers 2020-current. Admitted Ph.D. at Columbia University
Yilan Jiang 2021-2022. Now Ph.D. at UIUC
Jeffrey Tao 2021-current
Miles Huang 2022-current
Deema Alnuhait 2022-current
Shelly Cheng 2022-2022. Now at Splunk.
Joseph Rebagliati 2022-current
Yuhao Dong 2022
Qiran Li 2022
Lindsey Weiskopf 2022
Emily Ann Ye 2022

MS Advising Alexander Yao 2021-2022. Now at [Airtable](#).
Thesis: Towards Practical Provenance-based Query Explanation Engines

Sughosh V Kaushik 2021-2022. Now at [Eightfold.ai](#)
Thesis: Lineage Query Support in Columnar Database Engines

MS Thesis Committee Lucy X Wang.
Thesis: Modeling and Predicting the Dynamics of Clicks from Social Media

Undergrad Research Advising Hamed Nilforoshan 2016-2019. Ph.D. at Stanford.
Rahul Khanna 2016-2017. M.S. at USC.
Kevin Lin 2016-2018. AI2, Ph.D. at UC Berkeley.
Lily Xiu 2017 Summer. M.S at CMU, Ph.D. at UC Berkeley
Ian Yiran Huang 2017-2019. Ph.D. at Stanford.
Robert Netzorg 2017-2019. Ph.D. at UC Berkeley.
Lauren Arnett 2017-2019
Rodolfo Raimundo 2017-2018
Viraj Rai 2018
Sagar Lal 2018. M.S. at Columbia.
Maneet Khaira 2018
Amita Shukla 2018-2019
Condor Shou 2018-2019
Jacob Fisher 2019-current
Jeffrey Huang 2020
Gitika Bose 2020
Tejit Pabari 2020-2021
Alex Gajewski 2020-2021. Founded [metaphor.ai](#)
Alexander Yao 2021-2022
Vivek Anand Yanamadula 2021-2021
Hazel Zhu 2022-2022
Robert Ward 2022
Iman Anwarzai 2022
Ramya Subramanian 2023

High School Alex Studer 2017. Accepted to M.I.T
Joey O Connor 2016

INVITED TALKS

- 2023 Systems for Human Data Interaction *Keynote, North East Database Day*
- 2022 Systems for Human Data Interaction *Cornell University*
 Systems for Human Data Interaction *Sigma Computing*
 Panelist for New Research Symposium at SIGMOD
 Provenance and Data Markets *S.E.C. Quant Seminar*
 Systems for Human Data Interaction *Megagon Labs*
- 2021 Systems for Human Data Interaction *Parsons, Transdisciplinary Design Seminar*
 Systems for Human Data Interaction *Keynote at SEAS DATA Workshop at VLDB*
 Panelist for Interactive Querying and Visualization for Large Data Roundtable at VLDB
 Panelist for Data Governance and Provenance Roundtable at VLDB
 Systems for Human Data Interaction *CMU, Vaccination DB Talk Series*
- 2020 Systems for Human Data Interaction *Tufts University*
 Systems for Human Data Interaction *University of Chicago*
 Systems for Human Data Interaction *UCSD*
 Systems for Human Data Interaction *UC Berkeley*
- 2019 Systems for Human Data Interaction *Tsinghua University, Beijing, China*
 Systems for Human Data Interaction *China Big Data Workshop, Beijing, China*
 Tutorial: Towards Democratizing Relational Data Visualization *VLDB Summer School, Renmin University, Beijing, China*
 Precision Interfaces *SIGMOD in Amsterdam, Netherlands*
 Tutorial: Towards Democratizing Relational Data Visualization *SIGMOD in Amsterdam, Netherlands*
 Crazier Innovations in Databases \boxtimes Reinforcement-learning Research *CIDR in Monterey, CA*
 Human Data Interaction in the WuLab *UCLA*
- 2018 Human Data Interaction in the WuLab *Columbia Lamont Earth Institute*
 Human in the loop data analysis panel *NYDBDay, NY*
 Closing the loop on data analysis *NYDBDay, NY*
 Closing the loop on data analysis *UMass Amherst*
 Closing the loop on data analysis *BlueCore, NY*
 At a Glance: Approximate Entropy as a Measure of Line Chart Visualization Complexity
InfoVIS, Berlin, Germany
 Closing the loop on data analysis *Hasso Platner Institute, Potsdam, Germany*
 Closing the loop on data analysis *Data Science Institute, Columbia University*
 Closing the loop on data analysis *Harvard University*
 Closing the loop on data analysis *MIT*
 Closing the loop on data analysis *Google Research NYC*

- 2017 Closing the loop on data analysis *AT&T Research NYC*
Closing the loop on data analysis *University of Fribourg, Switzerland*
Closing the loop on data analysis *Dagstuhl, Germany*
Closing the loop on data analysis *UIUC*
Closing the loop on data analysis *U. Chicago*
Closing the loop on data analysis *U of Wisconsin-Madison*
Databases and Data Visualization *2Sigma, NY*
ICDE Panel: The Case for Small Data *ICDE in San Diego, CA*
CIDR: Chat-oriented Innovations in Database Research *CIDR in Santa Cruz, CA*
Combining Design and Performance in a Data Visualization Management System *CIDR in Santa Cruz, CA*
- 2016 Closing the loop on data analysis *Foundations of Data Science, DSI, Columbia University*
Closing the loop on data analysis *Sense, Collect & Move, DSI, Columbia University*
Closing the loop on data analysis *Brown University*
Towards Perception-aware Interactive Data Visualization Systems *DSIA in Chicago, IL*
Closing the loop on data analysis *Egelston Seminar at Columbia University*
Closing the loop on data analysis *IGERT Seminar at Columbia University*
PFunk-H: Approximate Query Processing using Perceptual Models *HILDA at SIGMOD, SF, CA*
Closing the loop on data analysis *IBM in Yorktown, NY*
Provenance in Big Data *MIT Big Data Workshop, Boston, MA*
Closing the loop on data analysis *Stats at Columbia University*
Closing the loop on data analysis *Rutgers University*
- 2015 Closing the loop on data analysis *IGERT Seminar at Columbia University*
VLDB Panel: Designing for Interaction: Broadening our View of Working with Data *VLDB in Hawaii*

COLUMBIA SERVICE

- 2022 Data Science Institute Co-Chair Center for Data, Media & Society
CS Chair for CS-Journalism Dual Degree
Computer Science Graduate Admissions Committee
- 2021 On parental leave
Data Science Institute Co-Chair Center for Data, Media & Society
- 2020 Data Science Institute Co-Chair Center for Data, Media & Society
CS Chair for CS-Journalism Dual Degree
Faculty mentor for Columbia Virtual Campus
DSI Data Art Contest Co-organizer
Computer Science Graduate Admissions Committee
- 2019 Data Science Institute Co-Chair Center for Data, Media & Society
CS Chair for CS-Journalism Dual Degree
DSI Data Art Contest Co-organizer
Computer Science Faculty Recruiting Committee
Computer Science Graduate Admissions Committee
- 2018 Data Science Institute Co-Chair Center for Data, Media & Society
Data Science Institute Masters Curriculum
Data Science Institute Center Committee: Center for Computing Systems for Data-Driven Science
CS Chair for CS-Journalism Dual Degree
Computer Science Graduate Admissions Committee
- 2017 Data Science Institute Co-Chair Center for Data, Media & Society
Data Science Institute Center Committee: Center for Computing Systems for Data-Driven Science
CS Chair for CS-Journalism Dual Degree
Computer Science Graduate Admissions Committee
- 2016 Data Science Institute Co-Chair Center for New Media
CS Chair for CS-Journalism Dual Degree
Computer Science Graduate Admissions Committee
- 2015 Computer Science Graduate Admissions Committee

COMMUNITY SERVICE

- 2024 CIDR PC
- 2023 SIGMOD PC
- SOCC PC
- DEEM PC
- HILDA PC & Mentor
- ICDE Workshop Co-Chair
- 2022 VLDB Associate Editor
- SIGMOD Demo PC
- HILDA PC & Mentor
- 2021 On parental leave
- VLDB Associate Editor
- SOCC PC
- NSF IIS Panelist
- 2020 VLDB PC
- Department of Energy Early Career PC
- 2019 NY DB Day Workshop Co-Chair
- SIGMOD PC
- SIGMOD Student Research Competition Co-chair
- NSF IIS Panel Reviewer
- 2018 ICDE PC
- SIGMOD PC
- HILDA Co-chair
- SIGMOD New Researcher Symposium Co-chair
- SIGMOD Publicity Co-chair
- Dagstuhl 17461: Connecting Visualization and Data Management Research
- 2017 ICDE Area Chair
- WWW PC
- SIGMOD Demo PC
- SIGMOD PC
- VLDB PC
- HILDA PC
- SSDBM PC
- HCOMP PC
- CHI Reviewer
- NSF IIS Panelist
- 2016 InfoVis Reviewer
- HILDA PC
- CHI Reviewer
- InfoVIS Reviewer
- NEDBDay Co-Chair
- SIGMOD travel award Co-chair
- 2015 SIGMOD travel award Co-chair
- 2014 DATA4U PC

Glossary

SIGMOD	ACM SIGMOD/PODS International Conference on Management of Data
VLDB	International Conference on Very Large Data Bases
ICDE	IEEE Conference on Data Engineering
CIDR	Conference on Innovative Data Systems Research
CHI	ACM CHI Conference on Human Factors in Computing Systems
WWW	The Web Conference
SOCC	ACM Symposium on Cloud Computing
SSDBM	International Conference on Scientific and Statistical Database Management
HCOMP	AAAI Conference on Human Computation and Crowdsourcing
HILDA	Human-in-the-Loop Workshop at SIGMOD
InfoVIS	IEEE Information Visualization conference
NEDBDay	New England Database Day
NYDBDay	New York Database Day
DATA4U	Workshop on Bringing the Value of “Big Data” to Users