

RESEARCH INTERESTS

I am broadly interested in data management systems that extend data analysis capabilities to non-expert users. Relevant fields include core database optimization, data provenance, and interface design.

EDUCATION

- Expected 2014 **Massachusetts Institute of Technology**, Cambridge, MA
Ph.D., Electrical Engineering and Computer Science
Advisor: Samuel Madden
Dissertation: Implementation and Applications of High Performance Provenance Systems for Data Analysis
- 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

- 2015– **Columbia University, NY, NY**
Assistant Professor – Computer Science
- 2015 **UC Berkeley, Berkeley, CA**
Visiting Scholar – AMPLab
- 2008–2014 **Massachusetts Institute of Technology, Cambridge, MA**
Ph.D. Student – CSAIL

CURRENT PROJECTS

Data Visualization Management Systems

A Data Visualization Management System (DVMS) integrates visualizations and databases, by compiling a declarative visualization language into an end-to-end relational operator pipeline that renders the visualization and is amenable to database-style optimizations. Thus the DVMS can be both expressive via the visualization language, and performant by leveraging traditional and visualization-specific optimizations to scale interactive visualizations to massive datasets.

<http://eugenewu.net/dvms.html>

Perceptual Functions for Data Visualization

Increasing data sizes has made it more difficult to build highly responsive interactive visualization tools due to the enormous quantity of input data and results that must be computed. Sampling-based approximation query processing is a promising research direction however over and under-sampling can easily lead to wasted resources or incorrect visualizations. We are modeling human perceptual limitations and using those models to automatically help visualization systems generate approximate but perceptually accurate visualizations.

<http://perceptvis.github.io/>

Data Cleaning for Machine Learning

Data cleaning is fundamentally challenging because there does not exist a pre-existing notion of correctness for the cleaned data. In addition, it is unclear whether data cleaning improves the downstream applications. For example, it is possible that data cleaning can make machine learning models worse than no cleaning at all. We are exploring semi and fully-automated data cleaning techniques for machine learning applications

<https://activeclean.github.io>

Query Explanation

Instead of explaining and fixing data using data, query explanation seeks to both explain and repair incorrect data values by using the actual queries that modified the database.

<http://eugenewu.net/qfix.html>

Data Cleaning Systems

Analysts report spending upwards of 80% of their time on problems in data cleaning including extraction, formatting, handling missing values, and entity resolution. How can knowing the application you want to actually run help speed up the cleaning process?

PAST PROJECTS AND JOBS

MIT Big Data Challenge

I developed and ran MIT's largest Big Data prediction and visualization challenge.

<http://bigdatachallenge.csail.mit.edu>

"Why" Analysis of SQL Aggregate Queries

I designed and implemented an analysis framework to explain outliers in the results of aggregation queries by constructing predicates on the input data. I formalized the concept of predicate influence and identified several operator properties to enable more efficient search algorithms on common statistical aggregates.

Efficient, Low Overhead Provenance

I designed, prototyped, and evaluated a low overhead provenance system for large-scale scientific workflow applications that process gigabytes of data per second.

Query Processing with Humans

This project pioneered the use of human computation platforms such as Mechanical Turk within a database query execution engine.

Index and Partitioning Techniques

I investigated the application of indexing and partitioning techniques for time-varying and skewed query workloads. Shinobi incrementally re-partitions and indexes database tables based on recent query access patterns. Our subsequent No Bits Left Behind paper proposed the use of unused space in B-tree index pages as a cache for heavily accessed tuples. This could improve the performance of skewed query workloads such as Wikipedia's access patterns by up to three orders of magnitude.

Trajectory Optimized Storage

I implemented the core storage system for TrajStore, a high performance data management system for storing and querying vehicle trajectory data by location and time. The system incrementally optimizes the storage layout as the query workload changes over time.

2007-2008 **Google Inc., Mountain View, CA**

Intern – Data Management Research

I worked in Alon Halevy's data management group on the WebTables project to mine the Google web corpus for tabular data. I developed the table extraction pipeline and extracted more than 125 million tables. In addition, I built a table search engine that lets users query over the structured data and automatically visualize attributes in graphs or maps.

Summer 2006 **Yahoo!, Santa Clara, CA**

Engineering Intern

I explored efficient implementations of RDF stores for an internal project.

Summer 2005 **Microsoft Inc., Redmond, WA**

Engineering Intern

I worked on efficient deep cloning and other internal features in Exchange Server

Spring 2005 **IBM Extreme Blue., Almaden, CA**

Engineering Intern

I developed a new software patch service for DB2 for z/OS team that reduced patch application times from the order of months to minutes.

2004–2006 **UC Berkeley, Berkeley, CA**

Undergraduate Researcher – Computer Science Department

High Performance Stream Processing

I designed and implemented one of the first high performance complex event processing systems for detecting high level events (e.g., shoplifting occurred) from streams of raw sensor events (e.g., RFID tag XXX detected). Our results were published at SIGMOD, the premier database conference.

The HiFi Project

I implemented the RFID reader interface for extracting raw events from early RFID readers and the interactive dashboard for the VLDB demonstration. HiFi is a research project around cascading stream architectures for large-scale geo-distributed receptor-based networks.

GRANTS

- 2016 ACM SIGMOD Conference 2016: Student Activities and Travel Support
- IIS: Medium: Collaborative Research: Composing Interactive Data Visualizations
- Columbia Alliance: Perceptual Functions for Faster Interactive Visualizations
- REU: Development of Graphical Perception as a Service
- 2015 IIS: Small: Collaborative Research: Towards Interactive Data Visualization Management Systems

SERVICE

- 2017 ICDE Area Chair
WWW PC
SIGMOD Demo PC
SIGMOD PC
VLDB PC
HILDA PC
SSDBM PC
HCOMP PC
- 2016 InfoVis Reviewer
HILDA PC
NEDBDay Co-Chair
SIGMOD travel award committee
- 2015 SIGMOD travel award committee
- 2014 DATA4U PC

TEACHING EXPERIENCE

- Spring 2016 *Instructor, Interactive Data Exploration Systems*
<http://columbiaviz.github.io>
- Instructor, Computing Systems for Data Science*
<http://w4121.github.io>
- Fall 2016 *Instructor, Introduction to Databases*
- Spring 2016 *Instructor, Big Data Systems*
- Fall 2015 *Instructor, Introduction to Databases*
- Fall 2013 *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>
- Spring 2012 *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.
- Fall 2010 *Teaching Assistant, Database Systems (MIT 6.830)*
 I assisted in writing and grading the assignments and projects.
- Summer 2010 *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.
- Spring 2010 *Instructor, Introduction to Java Course (MIT 6.S092)*
- Spring 2011 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>
- Fall 2006 *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.

PERSONAL

I love drawing and designing T-shirts and posters. I have created over 20 designs that have been printed and my shirts have been worn by thousands of people. The following link lists some of my designs.

<http://eugenewu.net/gallery.html>

PUBLICATIONS

- [1] D. Alabi and E. Wu. Pfunk-h: Approximate query processing using perceptual models. *HILDA*, 2016.
- [2] L. Battle, E. Benson, A. Parameswaran, and E. Wu. Indexing cost sensitive prediction. *arXiv preprint arXiv:1408.4072*, 2014.
- [3] A. Bhardwaj, A. Deshpande, A. J. Elmore, D. Karger, S. Madden, A. Parameswaran, H. Subramanyam, E. Wu, and R. Zhang. Collaborative data analytics with datahub. *Proceedings of the VLDB Endowment*, 8(12):1916–1919, 2015.
- [4] A. A. Bhattacharya, D. Hong, D. Culler, J. Ortiz, K. Whitehouse, and E. Wu. Automated metadata construction to support portable building applications. In *Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments*. ACM, 2015.
- [5] M. J. Cafarella, A. Halevy, D. Z. Wang, E. Wu, and Y. Zhang. Webtables: exploring the power of tables on the web. *Proceedings of the VLDB Endowment*, 1(1):538–549, 2008.
- [6] M. J. Cafarella, A. Y. Halevy, Y. Zhang, D. Z. Wang, and E. Wu. Uncovering the relational web. *under review*, 2008.
- [7] A. Cheung, L. Ravindranath, E. Wu, S. Madden, and H. Balakrishnan. Mobile applications need targeted micro-updates. In *Proceedings of the 4th Asia-Pacific Workshop on Systems*, page 8. ACM, 2013.
- [8] O. Cooper, A. Edakkunni, M. J. Franklin, W. Hong, S. R. Jeffery, S. Krishnamurthy, F. Reiss, S. Rizvi, and E. Wu. Hifi: A unified architecture for high fan-in systems. *PROCEEDINGS OF THE INTERNATIONAL CONFERENCE ON VERY LARGE DATA BASES*, pages 1357–1360, 2004.
- [9] P. Cudre-Mauroux, E. Wu, and S. Madden. The case for rodentstore, an adaptive, declarative storage system. *arXiv preprint arXiv:0909.1779*, 2009.
- [10] P. Cudre-Mauroux, E. Wu, and S. Madden. Trajstore: An adaptive storage system for very large trajectory data sets. In *Data Engineering (ICDE), 2010 IEEE 26th International Conference on*, pages 109–120. IEEE, 2010.
- [11] C. Curino, E. P. Jones, R. A. Popa, N. Malviya, E. Wu, S. Madden, H. Balakrishnan, and N. Zeldovich. Relational cloud: A database-as-a-service for the cloud. 2011.
- [12] M. J. Franklin, S. R. Jeffery, S. Krishnamurthy, F. Reiss, S. Rizvi, E. Wu, O. Cooper, A. Edakkunni, and W. Hong. *Design considerations for high fan-in systems: The HiFi approach*, volume 5. CIDR, 2005.
- [13] M. N. Garofalakis, K. P. Brown, M. J. Franklin, J. M. Hellerstein, D. Z. Wang, E. Michelakis, L. Tancau, E. Wu, S. R. Jeffery, and R. Aipperspach. Probabilistic data management for pervasive computing: The data furnace project. *IEEE Data Eng. Bull.*, 29(1):57–63, 2006.
- [14] D. Gyllstrom, E. Wu, H.-J. Chae, Y. Diao, P. Stahlberg, and G. Anderson. Sase: Complex event processing over streams. *arXiv preprint cs/0612128*, 2006.
- [15] D. Haas, S. Krishnan, J. Wang, M. J. Franklin, and E. Wu. Wisteria: Nurturing scalable data cleaning infrastructure. *Proceedings of the VLDB Endowment*, 8(12):2004–2007, 2015.
- [16] D. Haas, J. Wang, E. Wu, and M. J. Franklin. Clamshell: speeding up crowds for low-latency data labeling. *VLDB*, 2015.
- [17] A. Jindal, P. Rawlani, E. Wu, S. Madden, A. Deshpande, and M. Stonebraker. Vertexica: your relational friend for graph analytics! *Proceedings of the VLDB Endowment*, 7(13):1669–1672, 2014.
- [18] S. Krishnan, M. J. Franklin, K. Goldberg, J. Wang, and E. Wu. Activeclean: An interactive data cleaning framework for modern machine learning. In *Proceedings of the 2016 International Conference on Management of Data*, pages 2117–2120. ACM, 2016.
- [19] S. Krishnan, D. Haas, M. J. Franklin, and E. Wu. Towards reliable interactive data cleaning: A user survey and recommendations. In *HILDA*, 2016.

- [20] S. Krishnan, J. Wang, M. J. Franklin, K. Goldberg, T. Kraska, T. Milo, and E. Wu. Sampleclean: Fast and reliable analytics on dirty data. 2015.
- [21] S. Krishnan, J. Wang, E. Wu, M. J. Franklin, and K. Goldberg. Activeclean: interactive data cleaning for statistical modeling. *Proceedings of the VLDB Endowment*, 9(12):948–959, 2016.
- [22] S. Madden, E. Wu, S. Madden, Y. Zhang, E. Jones, C. Curino, et al. Relational cloud: The case for a database service. 2010.
- [23] A. Marcus, E. Wu, D. Karger, S. Madden, and R. Miller. Human-powered sorts and joins. *Proceedings of the VLDB Endowment*, 5(1):13–24, 2011.
- [24] A. Marcus, E. Wu, D. R. Karger, S. Madden, and R. C. Miller. Demonstration of quirk: a query processor for humanoperators. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*, pages 1315–1318. ACM, 2011.
- [25] A. Marcus, E. Wu, D. R. Karger, S. Madden, and R. C. Miller. Platform considerations in human computation. *Workshop on crowdsourcing and human computation*, 2011.
- [26] A. Marcus, E. Wu, D. R. Karger, S. R. Madden, R. C. Miller, et al. Crowdsourced databases: Query processing with people. 2011.
- [27] A. Marcus, E. Wu, and S. Madden. Data in context: Aiding news consumers while taming dataspace. *DBCrowd 2013*, 47, 2013.
- [28] H. Nilforoshan, J. Sands, K. Lin, R. Khanna, and E. Wu. Dialectic: Enhancing text input fields with automatic feedback to improve social content writing quality. *arXiv preprint arXiv:1701.06718*, 2017.
- [29] L. Sun, M. J. Franklin, J. Wang, and E. Wu. Skipping-oriented partitioning for columnar layouts. *Proceedings of the VLDB Endowment*, 10(4):421–432, 2016.
- [30] X. Wang, A. Meliou, and E. Wu. Qfix: Demonstrating error diagnosis in query histories. *SIGMOD*, 2016.
- [31] X. Wang, A. Meliou, and E. Wu. Qfix: Diagnosing errors through query histories. *SIGMOD*, 2017.
- [32] E. Wu. Research statement and agenda.
- [33] E. Wu. *Shinobi: Insert-aware partitioning and indexing techniques for skewed database workloads*. PhD thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, 2010.
- [34] E. Wu. Data visualization management systems. In *CIDR*, 2015.
- [35] E. Wu, L. Battle, and S. R. Madden. The case for data visualization management systems: Vision paper. *Proceedings of the VLDB Endowment*, 7(10):903–906, 2014.
- [36] E. Wu, P. Cudre-Mauroux, and S. Madden. Demonstration of the trajstore system. *Proceedings of the VLDB Endowment*, 2(2):1554–1557, 2009.
- [37] E. Wu, C. Curino, and S. Madden. No bits left behind. 2011.
- [38] E. Wu, Y. Diao, and S. Rizvi. High-performance complex event processing over streams. In *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*, pages 407–418. ACM, 2006.
- [39] E. Wu et al. *Explaining data in visual analytic systems*. PhD thesis, Massachusetts Institute of Technology, 2015.
- [40] E. Wu, L. Jiang, L. Xu, and A. Nandi. Graphical perception in animated bar charts. *arXiv preprint arXiv:1604.00080*, 2016.
- [41] E. Wu and S. Madden. Partitioning techniques for fine-grained indexing. In *Data Engineering (ICDE), 2011 IEEE 27th International Conference on*, pages 1127–1138. IEEE, 2011.
- [42] E. Wu and S. Madden. Scorpion: Explaining away outliers in aggregate queries. *VLDB*, 2013.
- [43] E. Wu, S. Madden, and M. Stonebraker. A demonstration of dbwipes: clean as you query. *Proceedings of the VLDB Endowment*, 5(12):1894–1897, 2012.

- [44] E. Wu, S. Madden, and M. Stonebraker. Subzero: a fine-grained lineage system for scientific databases. In *Data Engineering (ICDE), 2013 IEEE 29th International Conference on*, pages 865–876. IEEE, 2013.
- [45] Y. Wu, J. M. Hellerstein, and E. Wu. A devil-ish approach to inconsistency in interactive visualizations. *HILDA*, 2016.