Curriculum Vitae — Oliver Kennedy

Department of Computer Science and Engineering
University at Buffalo, SUNY
338 Davis Hall
Buffalo, NY 14260

63 Bassett Road
Amherst, NY 14221

Phone: +1 (716) 645-1597
Email: okennedy@buffalo.edu
Web: https://odin.cse.buffalo.edu/

Education

PhD, Cornell University, Ithaca, NY — 2011
Department of Computer Science
Thesis: *"Watch out for... What?: Monitoring and Uncertainty in Scientific Computing"*
Advisor: Christoph Koch

MS, Cornell University, Ithaca, NY — 2008
Department of Computer Science

BS (with Honors), New York University, New York, NY — 2005
Department of Computer Science

BE, Stevens Institute of Technology, Hoboken, NJ — 2005
Department of Computer Engineering

Employment History

- Associate Professor, Dept. of Comp. Sci. & Eng., SUNY Buffalo — Sept. 2012 – present
- Postdoctoral Fellow, D.A.T.A. Lab, EPFL, Switzerland — June 2011 – Aug. 2012
- Intern, Microsoft Research, Redmond, WA — May 2010 – Aug. 2010
- Intern, Yahoo! Labs, Sunnyvale CA — May 2009 – Aug. 2009
- Teaching/Research Assistant, Cornell University — Sept. 2005 – May 2011
- Programmer, Humanities Computing Group, New York University — Sept. 2003 – Aug. 2005
- Programmer, Parallel and Dist. Sys. Group, New York University — Jan. 2001 – May 2003

Honors

The pVLDB 2012 paper titled "DBToaster: Higher-order Delta Processing for Dynamic, Frequently Fresh Views" was accepted to the "Best of VLDB 2012" issue of the VLDB Journal.

Oliver received a UB CSE Outstanding Funding Award for attracting over $3m of funding to the University at Buffalo in 2016.

The SIGMOD 2017 paper titled "Beta Probabilistic Databases: A Scalable Approach to Belief Updating and Parameter Learning", authored by Oliver's student Niccolo Meneghetti, was accepted as a Best-Of-SIGMOD extended article in the ACM Transactions on Database Systems.

Oliver received the UB School of Engineering and Applied Sciences Early Career Outstanding Teacher of the Year Award in 2017

Oliver received the NSF's CAREER award in 2018

Oliver received the UB Exceptional Scholar: Young Investigator Award in 2018

Oliver was recognized as a VLDB Distinguished Reviewer for his service on the PC of VLDB 2018

Oliver's outreach efforts were featured on Channel 4 news Buffalo[1].

The SIGMOD 2019 paper titled "Uncertainty Annotated Databases - A Lightweight Approach for Approximating Certain Answers" received the 2020 SIGMOD Reproducability Award

Oliver was recognized as a Distinguished PC member for SIGMOD 2020

Oliver received the CSE Departmental Senior Teaching Award in 2020

Oliver received the CSE Departmental Senior Research Award in 2021

Oliver received the School of Engineering and Applied Sciences Senior Teaching Award in 2021

Oliver received the IEEE Region 1 Technological Innovation Award

Oliver was recognized as a Distinguished Reviewer for pVLDB Vol 16 (VLDB 2023)

Oliver was recognized as a Distinguished Associate Editor for SIGMOD 2025

Oliver was recognized as a Distinguished Reviewer for pVLDB Vol 18 (VLDB 2025)

Students Advised

[1] https://www.wivb.com/news/local-news/students-in-buffalo-learning-how-to-build-advanced-technology-this-summer/1323587691/

PhD Students Advised (8 graduated, 4 current)

Niccolò Meneghetti (Graduated in 2016; Jointly advised with Jan Chomicki), Ying Yang (Graduated in 2017), Gökhan Kul (Graduated in 2018; Jointly advised with Shambhu Upadhyaya), Ting Xie (Graduated in 2019), William Spoth (Graduated in 2022), Carl Nuessle (Graduated in 2024; Jointly advised with Luke Ziarek), Poonam Kumari (Graduated in 2024), Aaron Huber (Graduated in 2025; Jointly advised with Atri Rudra, Zhuoyue Zhao, and Jaroslaw Zola), Darshana Balakrishnan (Current Student; Jointly advised with Luke Ziarek), Nick Brown (Current Student), Pratik Pokharel (Current Student), Victoria Dib (Current Student)

MS Students Advised (11 graduated, 0 current)

Ankur Upadhyay (Graduated in 2014; Jointly advised with Lukasz Ziarek), Nikhil Kataria (Graduated in 2014), Palaniappan Meiyappan (Graduated in 2014), Sumit Agarwal (Graduated in 2014; Jointly advised with Lukasz Ziarek), Saket Adusumilli (Graduated in 2015), Vinayak Karuppasamy (Graduated in 2015), Arindam Nandi (Graduated in 2016), Anand Sankar Bhagavandas (Graduated in 2017), Shivang Aggarwal (Graduated in 2017), Sneha Krishnamurthy (Graduated in 2017), Gourab Mitra (Graduated in 2018; Jointly advised with Luke Ziarek)

BS Students Advised (5 graduated, 0 current)

Daniel Bellinger (Graduated in 2014; Jointly advised with Lukasz Ziarek), Thomas Mitchell (Graduated in 2015; Jointly advised with Varun Chandola and Shambhu Upadhyaya), Grant Wrazen (Graduated in 2017; Jointly advised with Luke Ziarek and Geoff Challen), Hank Lin (Graduated in 2017; Jointly advised with Luke Ziarek), Olivia Alphonce (Graduated in 2017)

Thesis Committees (Excluding joint advisees; 2 graduated)

Duc Thanh Luong (May 2018), Sai Vikneshwar (May 2020)

Professional Activities

Chair Positions

- Reproducibility co-chair: Proceedings of the VLDB Endowment 2018 – 2022
- Scholarship co-chair: Symposium on Cloud Computing 2019
- Technical co-chair: International Conference on Very Large Data Bases 2020
- Treasurer: The ACM SIGMOD/PODS Conference 2022

Associate Editor Positions

- Conference: The ACM SIGMOD/PODS Conference (SIGMOD) 2024 – 2025

Program Comittee Member

- Workshop: The VLDB PhD Workshop (VLDB-PhD) 2013
- Demo: VLDB Demo Track (VLDB-Demo) 2016 – 2017

- Journal: Proceedings of the VLDB Endowment (pVLDB) 2017 – 2018, 2020 – 2025
- Conference: The ACM SIGMOD/PODS Conference (SIGMOD) 2015 – 2017, 2019 – 2023
- Conference: The International Conference on Data Engineering (ICDE) 2019
- Workshop: Provenance Week (PWEEK) 2016
- Workshop: Workshop on Human-In-the-Loop Data Analytics (HILDA) 2016 – 2020
- Conference: Symposium on Cloud Computing (SoCC) 2017

Reviewer

- Journal: The VLDB Journal (VLDBJ) 2013, 2017 – 2018, 2021 – 2022
- Journal: IEEE Transactions on Knowledge and Data Engineering (TKDE) 2013 – 2014
- Journal: ACM Transactions on Database Systems (TODS) 2015, 2018
- Journal: Computer Science Education (CSE) 2015
- Conference: The ACM SIGMOD/PODS Conference (SIGMOD) 2024 – 2025
- Journal: Transactions on Information Technology (TOIT) 2016
- Journal: Knowledge and Information Systems (KAIS) 2017
- Conference: The International Conference on Database Theory (ICDT) 2017

Service

- WebDB Session Chair 2013
- NSF Panel Member 2014 – 2016, 2018, 2020 – 2022, 2024 – 2025
- SIGMOD Reproducibility Committee 2015 – 2017

Professional Memberships

- ACM 2012 – present
- ACM-SIGMOD 2010 – present
- ACM-Computer Science Teachers Assoc. 2012 – present
  - WNY Chapter Secretary Sept. 2013 – Sept. 2015
- IEEE 2009 – present

Volunteer Work

- LIBERTY Partnerships 2013 – present
  - Summer Mentor 2013 – 2014
  - 'I Can' Program Volunteer 2015 – 2017
- Coder Dojo Mentor 2013 – 2014
- Science is Elementary Class Leader 2014 – 2016
- CSExplore 2019, 2025
- CSE Welcome Weekend 2019, 2021

Departmental Service

- Dept. of Comp. Sci. and Eng., University at Buffalo, SUNY

- Graduate Admissions Committee 2012 – 2017
- Graduate Advisory Committee 2012 – 2018, 2022 – 2025
  - Graduate Co-Director 2022 – 2025
- Undergraduate Advisory Committee 2021
- Distinguished Speaker Committee 2017
- Faculty Search Committee 2012 – 2014
- Facilities Committee 2016
- Colloquium Committee 2012 – 2014
  - Chair 2014
- Grievances Committee 2015 – 2016
- Undergraduate Curriculum Development Subcommittee 2015 – 2016
- Teaching Effectiveness Committee 2019, 2022
- Preserving Academic Integrity 2021 – 2022
- Teaching Load Task Force 2021 – 2023
  - Co-Chair 2022 – 2023

Invited Talks

1. *Minnowbrook @ Syracuse* — Draupnir: A Database Engine for Declarative Compilers May 2025
2. *UC Santa Cruz* — Draupnir: A Database Engine for Declarative Compilers Apr. 2025
3. *UC Berkeley* — Draupnir: A Database Engine for Declarative Compilers Apr. 2025
4. *Institute for Defense Analytics: Center for Computing Sciences* — Draupnir: A Database Engine for Declarative Compilers Feb. 2025
5. *New York University* — Data Preparation with Vizier Apr. 2024
6. *University of Illinois: Chicago* — Principled management of notebook state in Vizier Apr. 2024
7. *University of Massachusetts: Dartmouth* — ASTral: A Declarative Compiler Compiler Mar. 2024
8. *Cornell University* — Microkernel Notebooks Feb. 2023
9. *UP-STAT 2022* — Panel: On the Multifaceted Impact of Artificial Intelligence in Healthcare: Past, Present, and Emerging Trends May 2022
10. *MARDA 2022* — Caveatting your data: Adding explainability to incomplete datasets Feb 2022
11. *University at Buffalo: Center for Global Health Equity* — How to Start Collaborating with CSE to Solve Global Health Problems May 2020
12. *Northeastern University* — Safe, Reusable Heuristic Data Transformation (through Caveats) Dec 2019
13. *Tableau, Inc.* — Safe, Reusable Heuristic Data Transformation (through Caveats) Dec 2019
14. *University of Massachusetts Amherst* — Safe, Reusable Heuristic Data Transformation (through Caveats) Dec 2019
15. *University of Waterloo* — Just-In-Time Data Structures May. 2019
16. *Cornell* — Don't Wrangle, Guess Instead (with Mimir) Jan. 2018
17. *Penn State U.* — Don't Wrangle, Guess Instead (with Mimir) Jan. 2018
18. *Harvard* — Just-In-Time Data Structures Dec. 2017
19. *University of Washington* — Don't Wrangle, Guess Instead (with Mimir) Dec. 2017
20. *Columbia* — Don't Wrangle, Guess Instead (with Mimir) Oct. 2017
21. *University of Michigan* — Don't Wrangle, Guess Instead (with Mimir) Oct. 2017

22. *Illinois Inst. Tech.* — Don't Wrangle, Guess Instead (with Mimir) Oct. 2017
23. *University of Chicago* — Don't Wrangle, Guess Instead (with Mimir) Oct. 2017
24. *Ohio State University* — Don't Wrangle, Guess Instead (with Mimir) Oct. 2017
25. *Cornell* — Just-In-Time Data Structures Sep. 2017
26. *ICDE* — Small Data (panel moderator) Apr. 2017
27. *Buffalo Database Seminar Meetup* — Post-Truth Databases (lightning talk) Apr. 2017
28. *Airbus* — Embracing Uncertainty with Mimir Oct. 2016
29. *RIT* — Pocket Data Oct. 2016
30. *LogicBlox* — Just-In-Time Data Structures Jun. 2016
31. *NYU* — Embracing Uncertainty with Mimir May. 2016
32. *HPE/Vertica* — Embracing Uncertainty with Mimir Mar. 2016
33. *Carnegie Mellon* — Embracing Uncertainty with Mimir Feb. 2016
34. *Buffalo District Board of Education* — LEGO NXT Professional Development Jan. 2015
35. *Ohio State University* — PocketData: What's in your pocket? Oct. 2015
36. *Oracle* — Just-In-Time Data Structures Jan. 2015
37. *Pivotal* — Just-In-Time Data Structures Jan. 2015
38. *Snowflake Computing Inc.* — Just-In-Time Data Structures Jan. 2015
39. *Western New York CSTA Fall Conference* — Queries with Twitter Tutorial Oct. 2014
40. *Buffalo Database Seminar Meetup* — Everything I Need to Know I Learned From PL Sept. 2014
41. *Google CS4HS, ACM-CSTA, and Buffalo State* — Mobilize Prime Workshop (with Lukasz Ziarek, Sarbani Banerjee) March. 2014
42. *Western New York CSTA Professional Development Workshops* — SQL Tutorial Jan. 2014
43. *Buffalo Database Seminar Meetup* — Emerging Trends in Database Research Jan. 2014
44. *Western New York CSTA 2013 CS4HS Workshop* — Python Workshop (with Lukasz Ziarek, Sarbani Banerjee) Jun. 2013
45. *Cornell University Database Colloquium* — Laasie: The Log as a Service Infrastructure Feb. 2013
46. *Western New York CSTA 2012 Fall Conference* — Python Workshop (with Lukasz Ziarek, Sarbani Banerjee) Oct. 2012

Courses Taught

Department of Computer Science and Engineering, SUNY Buffalo

- `CSE 4/562` : (planned) Database Systems — 100 students Spring 2026
- `CSE 350` : (planned) Advanced Data Structures — 8 students Fall 2025
- `CSE 501` : Grad Studies in CS — 82 students Fall 2024
- `CSE 250` : Data Structures — 94 students Fall 2024
- `CSE 410 [Proposed CSE 350]` : Advanced Data Structures — 10 students Spring 2024
- `CSE 250` : Data Structures — 122 students Fall 2023
- `CSE 501` : Grad Studies in CS — 24 students Fall 2022
- `CSE 250` : Data Structures — 152 students Fall 2022
- `CSE 250` : Data Structures — 266 students Fall 2021
- `CSE 715` : Data for ... Equitable Machine Learning — 7 students Spring 2021
- `CSE 4/562` : Database Systems — 62 students Spring 2021
- `CSE 611` : MS Project Development — 28 students Fall 2020

- `CSE 662` : Languages and Runtimes for Big Data — Fall 2019
- `CSE 4/562` : Database Systems — 59 students — Spring 2019
- `CSE 611` : MS Project Development — 46 students — Fall 2018
- `CSE 4/562` : Database Systems — 93 students — Spring 2018
- `CSE 662` : Languages and Runtimes for Big Data — 13 students — Fall 2017
- `CSE 562` : Graduate Database Systems — 54 students — Spring 2017
- `CSE 199` : Undergraduate CS Seminar Recitation — 21 students — Fall 2016
- `CSE 705` : Database Research Seminar — 7 students — Fall 2016
- `CSE 662` : Languages and Runtimes for Big Data — 10 students — Fall 2016
- `CSE 462` : Undergraduate Database Systems — 72 students — Spring 2016
- `CSE 662` : Languages and Runtimes for Big Data — 29 students — Fall 2015
- `CSE 562` : Graduate Database Systems — 79 students — Spring 2015
- `CSE 704` : Databases, Programming Languages and Datastructures — 21 students — Fall 2014
- `CSE 562` : Graduate Database Systems — 148 students — Spring 2014
- `CSE 704` : Streaming, Incremental and Online Data Processing — 18 students — Fall 2013
- `CSE 562` : Graduate Database Systems — 101 students — Spring 2013
- `CSE 704` : Web-Scale Data Management Systems Seminar — 15 students — Fall 2012

School of Computer and Communication Sciences, EPFL, Switzerland

- `CS-422` : Advanced Databases (as Teaching Assistant) — Spring 2011

Department of Computer Science, Cornell University

- `CS 4410` : Operating Systems (as Practicum Instructor) — Fall 2008
- `CS 414` : Operating Systems (as Practicum Instructor) — Spring 2006
- `CS 414` : Operating Systems (as Practicum Instructor) — Fall 2006

Grant Support

Pending Applications (2)

Title :	CSSI: Frameworks: Theseus: Semi-Homogeneous Data Integration for Related Survey Instruments
Agency :	NSF
Role :	PI
Start :	08/2025
End :	08/2030
Amount :	$1,420,140
Effort :	20%
Co-PIs :	Jeff Good, Andrew H. Talal, Marianthi Markatou, Katarzyna Kordas

Title :	MFAI: Arithmetic Circuits with Curated Uncertainty for Reliable Analyses, Training, and Estimation (ACCURATE)
Agency :	NSF
Role :	Co-PI
Start :	05/2025
End :	04/2028
Amount :	$1,420,140
Effort :	25%
Co-PIs :	Atri Rudra, Maria Rodriguez, Kenneth Joseph

Active Grants (2) — $732,922 Total; $466,460 By Effort

Title :	Collaborative Research: III: MEDIUM: U4U - Taming Uncertainty with Uncertainty-Annotated Databases
Agency :	NSF: CISE: IIS: MEDIUM
Role :	PI
Start :	08/15/2020
End :	08/14/2025
Amount :	$532,923 (Total Across Collaboration: $999,492)
Effort :	50%
Co-PI :	Atri Rudra
Peer Org. :	Illinois Inst. Tech. ($466,569)

Title :	STTR Phase 2: Curating Uncertainty for Reliable Exploitation and Collaboration (CURE-C)
Agency :	NASA
Role :	PI
Start :	04/01/2025
End :	03/15/2027
Amount :	$199,999 (Total Across Collaboration: $1,799,999)
Effort :	100%
Peer Org.s :	Xanalytix Systems ($0.00) Breadcrumb Analytics ($1,600,000)

Completed (15) — $5,900,782 Total; $2,643,069 By Effort

Title :	Physical Layout Optimization & Query Transformation for messy JSON collection
Agency :	Oracle University Relations
Role :	PI
Start :	05/2020
End :	05/2021
Amount :	$95,635
Effort :	100%

Title :	CAREER: Declarative Uncertainty
Agency :	NSF: CISE: IIS: III
Role :	PI
Start :	02/2018
End :	02/2024
Amount :	$542,276
Effort :	100%

Title :	Adaptive Schema and Elastic Query for No-SQL Data
Agency :	Oracle University Relations
Role :	PI
Start :	12/2017
End :	12/2018
Amount :	$96,684
Effort :	100%

Title :	CIF21 DIBBs: EI: Vizier, Streamlined Data Curation
Agency :	NSF: ACI: DIBBS
Role :	PI
Start :	01/2017
End :	06/2021
Amount :	$2,749,699
Effort :	33%
Co-PIs :	Boris Glavic, Juliana Freire
Supplements :	NSF-REU (amount included in total above)

Title :	CI-P: Supporting Pocket Scale Data Management Research
Agency :	NSF: CISE: IIS: CRI
Role :	PI
Start :	08/2016
End :	08/2017
Amount :	$100,000
Effort :	33%
Co-PIs :	Lukasz Ziarek, Geoffrey Challen

Title :	Curating Uncertainty and Reliable Exploitation (CURE)
Agency :	The US Naval Postgraduate School
Role :	Co-I
Start :	08/2016
End :	12/2017
Amount :	$265,357
Effort :	50%
Co-PI :	Moises Sudit

Title :	III: Small: Just in Time Datastructures
Agency :	NSF: CISE: IIS: III
Role :	PI
Start :	07/2016
End :	04/2019
Amount :	$502,274
Effort :	50%
Co-PI :	Lukasz Ziarek
Supplements :	NSF-REU (amount included in total above)

Title :	Intuitive Data Interpretation
Agency :	Oracle University Relations
Role :	PI
Start :	05/2016
End :	05/2017
Amount :	$89,187
Effort :	100%

Title :	SCC-PG: A Sustainable and Connected Community-Scale Food System to Empower Consumers, Farmers, and Retailers
Agency :	NSF: SCC
Role :	Co-PI
Start :	08/15/2021
End :	08/14/2023
Amount :	$150,000
Effort :	20%
Co-PIs :	Samina Raja, Sara Behdad, Debabrata Talukdar, Srirangaraj Setlur, Emmanuel Frimpong Boamah

Title :	Expressing Uncertainty Using the maybe System
Agency :	Google Research Awards
Role :	Co-PI
Start :	08/2015
End :	07/2016
Amount :	$38,656
Effort :	33%
Co-PIs :	Geoffrey Challen, Lukasz Ziarek

Title :	Intuitive Data Interpretation
Agency :	Oracle University Relations
Role :	PI
Start :	03/2015
End :	03/2016
Amount :	$90,455
Effort :	100%

Title :	TWC: Medium: Collaborative: Data is Social: Exploiting Data Relationships to Detect Insider Attacks
Agency :	NSF: CISE: SaTC
Role :	PI
Start :	09/2014
End :	08/2019
Amount :	$975,999 (Total Across Collaboration: $1,215,973)
Effort :	25%
Co-PIs :	Hung Ngo, Shambhu Upadhyaya, Varun Chandola
Peer Org. :	University of Michigan Ann Arbor ($239,974)
Supplements :	NSF-REU (amount included in total above)

Title :	Intuitive Data Interpretation
Agency :	Oracle University Relations
Role :	PI
Start :	03/2014
End :	03/2015
Amount :	$88,000
Effort :	50%
Co-PI :	Jan Chomicki

Title :	NSF Student Travel Grant for 2019 Symposium on Cloud Computing (SOCC)
Agency :	NSF: CISE: CNS: CSR
Role :	PI
Start :	10/01/2019
End :	09/30/2020
Amount :	$15,000
Effort :	100%

Title :	STTR Phase 1: Curating Uncertainty for Reliable Exploitation and Collaboration (CURE-C)
Agency :	NASA
Role :	PI
Start :	08/01/2023
End :	07/30/2024
Amount :	$149,560 (Total Across Collaboration: $149,560)
Effort :	50%
Peer Org. :	Xanalytix Systems ($0.00)

Publications and Artifacts

Students I advise or co-advise and I are in **bold**, my former advisor is underlined

Journal Publications (9)

M. Markatou, **O. Kennedy**, M. Brachmann, R. Mukhopadhyay, A. Dharia, A. H. Talal; Social determinants of health derived from people with opioid use disorder: Improving data collection, integration and use with cross-domain collaboration and reproducible, data-centric, notebook-style workflows — *Frontiers in Medicine (Frontiers in Medicine) 2023 (Volume 10, Number),* 19 pages, acceptance rate not available

S. Feng, B. Glavic, **O. Kennedy**; Efficient Approximation of Certain and Possible Answers for Ranking and Window Queries over Uncertain Data — *Proceedings of the VLDB Endowment (pVLDB) 2023,* 12 pages, acceptance rate 24.8%

O. Kennedy, B. Glavic, J. Freire, M. Brachmann; The Right Tool for the Job: Data-Centric Workflows in Vizier — *IEEE Data Engineering Bulletin (IEEE-DEB) 2022,* 13 pages, acceptance rate not available

T. Xie, V. Chandola, **O. Kennedy**; Query Log Compression for Workload Analytics — *Proceedings of the VLDB Endowment (pVLDB) 2019,* 12 pages, acceptance rate 18.91%

N. Meneghetti, **O. Kennedy**, W. Gatterbauer; Learning From Query-Answers: A Scalable Approach to Belief Updating and Parameter Learning — *ACM Transactions on Database Systems (TODS) 2018,* 40 pages, acceptance rate not available / Invited article extending a 'Best-of-SIGMOD' paper from SIGMOD 2017

G. Kul, **D. T. Luong**, **T. Xie**, V. Chandola, **O. Kennedy**, S. Upadhyaya; Similarity Metrics for SQL Query Clustering — *IEEE Transactions on Knowledge and Data Engineering (TKDE) 2018 (Volume 30, Number),* 15 pages, acceptance rate not available

Y. Yang, **N. Meneghetti**, R. Fehling, Z. H. Liu, D. Gawlick, **O. Kennedy**; Lenses: An On-Demand Approach to ETL — *Proceedings of the VLDB Endowment (pVLDB) 2015 (Volume 8, Number 12),* 12 pages, acceptance rate 21.3%

C. Koch, Y. Ahmad, **O. Kennedy**, M. Nicolic, A. Nötzli, D. Lupei, A. Shaikhana (*Authors listed in descending order of seniority*; DBToaster: Higher-order Delta Processing for Dynamic, Frequently Fresh Views — *The VLDB Journal (VLDBJ) 2014 (Volume 23, Number 2; Special issue on the best papers of VLDB 2012),* 25 pages, acceptance rate not available / Extended version of a conference paper, published in a special 'Best-of-VLDB' issue of VLDBJ

Y. Ahmad, **O. Kennedy**, C. Koch, M. Nicolic; DBToaster: Higher-order Delta Processing for Dynamic, Frequently Fresh Views — *Proceedings of the VLDB Endowment (pVLDB) 2012,* 12 pages, acceptance rate 20.3% / Invited to submit an extended version to a 'Best-of-VLDB' issue of VLDBJ

Conference Publications (20)

A. Huber, **O. Kennedy**, A. Rudra, Z. Zhao, S. Feng, B. Glavic; FastPDB: Towards Bag-Probabilistic Queries at Interactive Speeds — *The ACM SIGMOD/PODS Conference (SIGMOD)* 2025; Berlin, Germany; 13 pages, acceptance rate 24.8%

D. Balakrishnan, **C. Nuessle**, **O. Kennedy**, L. Ziarek; TreeToaster: Towards an IVM-Optimized Compiler — *The ACM SIGMOD/PODS Conference (SIGMOD)* 2021; Xi'an, Shaanxi, China [Online]; 12 pages, acceptance rate unknown; average rate for past conferences approximately 19.25%

S. Feng, B. Glavic, **A. Huber**, **O. Kennedy**; Efficient Uncertainty Tracking for Complex Queries with Attribute-level Bounds — *The ACM SIGMOD/PODS Conference (SIGMOD)* 2021; Xi'an, Shaanxi, China [Online]; 12 pages, acceptance rate unknown; average rate for past conferences approximately 19.25%

W. Spoth, **O. Kennedy**, Y. Lu, B. Hammerschmidt, Z. H. Liu; Reducing Ambiguity in Json Schema Discovery — *The ACM SIGMOD/PODS Conference (SIGMOD)* 2021; Xi'an, Shaanxi,

China [Online]; 12 pages, acceptance rate unknown; average rate for past conferences approximately 19.25%

P. Kumari, M. Brachmann, **O. Kennedy**, S. Feng, B. Glavic; DataSense: Display Agnostic Data Documentation — *The Conference on Innovative Data Systems Research (CIDR)* 2021; Online-Only; 1 page, acceptance rate unknown (less than 30% in previous years)

M. Brachmann, **W. Spoth**, **O. Kennedy**, B. Glavic, H. Mueller, S. Castelo, C. Bautista, J. Freire; Your notebook is not crumby enough, REPLace it — *The Conference on Innovative Data Systems Research (CIDR)* 2020; Amsterdam, NL; 14 pages, acceptance rate unknown (less than 30% in previous years)

D. Balakrishnan, L. Ziarek, **O. Kennedy**; Fluid Data Structures — *Proceedings for DBPL (DBPL)* 2019; Phoenix, AZ, USA; 15 pages, acceptance rate not available

S. Feng, **A. Huber**, B. Glavic, **O. Kennedy**; Uncertainty Annotated Databases - A Lightweight Approach for Approximating Certain Answers — *The ACM SIGMOD/PODS Conference (SIGMOD)* 2019; Amsterdam, NL; 18 pages, acceptance rate unknown; average rate for past conferences approximately 19.25%

N. Meneghetti, **O. Kennedy**, W. Gatterbauer; Beta Probabilistic Databases: A Scalable Approach to Belief Updating and Parameter Learning — *The ACM SIGMOD/PODS Conference (SIGMOD)* 2017; Chicago, IL, USA; 12 pages, acceptance rate 19.6% / Invited to submit an extended version as a 'Best-of-SIGMOD' paper to ACM-TODS

Y. Yang, **O. Kennedy**; Convergent Inference with Leaky Joins — *International Conference on Extending Database Technology (EDBT)* 2017; Venice, Italy; 12 pages, acceptance rate 21.8%

W. Spoth, B. S. Arab, E. S. Chan, D. Gawlick, A. Ghoneimy, B. Glavic, B. Hammerschmidt, **O. Kennedy**, S. Lee, Z. H. Liu, X. Niu, **Y. Yang** (*Authors listed in alphabetical order*); Adaptive Schema Databases — *The Conference on Innovative Data Systems Research (CIDR)* 2017; Santa Cruz, CA, USA; 6 pages, acceptance rate unknown (less than 30% in previous years)

O. Kennedy, L. Ziarek; Just in Time Datastructures — *The Conference on Innovative Data Systems Research (CIDR)* 2015; Monterey, CA, USA; 10 pages, acceptance rate unknown (less than 30% in previous years)

K. Jeon, S. Chandrashekhara, F. Shen, S. Mehra, **O. Kennedy**, S. Y. Ko (*Authors listed students first, then faculty*; PigOut: Making Multiple Hadoop Clusters Work Together — *The IEEE International Conference on Big Data (IEEE BigData)* 2014; Washington, DC, USA; 9 pages, acceptance rate 18.6%

J. Yang, E. Vee, S. Vassilvitskii, J. Tomlin, J. Shanmugasundaram, T. Anastasakos, **O. Kennedy**; Inventory Allocation for Online Graphical Display Advertising Using Multi-Objective Optimization — *The International Conference on Operations Research and Enterprise Systems (ICORES)* 2012; Vilamoura, Portugal; 12 pages, acceptance rate not available

O. Kennedy, S. Nath; Jigsaw: Efficient optimization over uncertain enterprise data — *The ACM SIGMOD/PODS Conference (SIGMOD)* 2011; Athens, Greece; 12 pages, acceptance rate 25.0%

O. Kennedy, Y. Ahmad, C. Koch; DBToaster: Agile Views for a Dynamic Data Management System — *The Conference on Innovative Data Systems Research (CIDR)* 2011; Monterey, CA, USA; 12 pages, acceptance rate unknown (less than 30% in previous years)

O. Kennedy, C. Koch; PIP: A Database System for Great and Small Expectations — *The International Conference on Data Engineering (ICDE)* 2010; Long Beach, CA, USA.; 12 pages, acceptance rate 21.0%

O. Kennedy, C. Koch, A. Demers; Dynamic approaches to in-network aggregation — *The International Conference on Data Engineering (ICDE)* 2009; Shanghai, China.; 4 pages, acceptance rate 17.0%

G. Shawver, **O. Kennedy**; The Face of Meaning — *The Joint International Conference of the Association for Computers and the Humanities and the Association for Literary and Linguistic Computing (ACH/ALLH)* 2005; Victoria, BC, Canada; 6 pages, acceptance rate not available

E. Lage-Otero, **O. Kennedy**; Second language comprehension in an online multimedia environment: Integrating theory and practice. — *The World Conference on E-Learning in Corporate, Government, Healthcare, and Higher Education (E-Learn)* 2004; Washington, DC, USA; 2 pages, acceptance rate not available

Workshop Publications (20)

P. Pokharel, **O. Kennedy**; Embedding The Context of Attributes in Longitudinal Study Schemas To Refine Their Semantic Match Candidates — *Ontario Database Day (OnDBD)* 2024; Waterloo, ON, Canada; 1 page, acceptance rate not available

N. Brown, **O. Kennedy**; On-Disk Datalog Engine using Semirings — *Ontario Database Day (OnDBD)* 2024; Waterloo, ON, Canada; 1 page, acceptance rate not available

D. Balakrishnan, **O. Kennedy**, L. Ziarek, J. Luong, H. Gildhoff, G. Saxena; Multiquery Optimization for Declarative Compilers — *Ontario Database Day (OnDBD)* 2024; Waterloo, ON, Canada; 1 page, acceptance rate not available

P. Pokharel, J. Lee, **O. Kennedy**, J. Good, M. Markatou, A. H. Talal, R. Mukhopadhyay; Drag, Drop, Merge: A Tool for Streamlining Integration of Longitudinal Survey Instruments — *Workshop on Human-In-the-Loop Data Analytics (HILDA)* 2024; Santiago, Chile; 6 pages, acceptance rate unknown; average rate for past conferences approximately 50.0%

O. Kennedy, B. Glavic, M. Brachmann; Overlay Spreadsheets — *Workshop on Human-In-the-Loop Data Analytics (HILDA)* 2023; Seattle, WA, USA; 7 pages, acceptance rate unknown; average rate for past conferences approximately 50.0%

N. Deo, B. Glavic, **O. Kennedy**; Runtime Provenance Refinement for Notebooks — *USENIX Workshop on Theory and Practice of Provenance (TaPP)* 2022; Philadelphia, NY; 4 pages, acceptance rate 58.82%

W. Spoth, **P. Kumari**, **O. Kennedy**, F. Nargesian; Loki: Streamlining Integration and Enrichment — *Workshop on Human-In-the-Loop Data Analytics (HILDA)* 2020; Online-Only (originally Portland, OR, USA); 4 pages, acceptance rate unknown; average rate for past conferences approximately 50.0%

C. Nuessle, **O. Kennedy**, L. Ziarek; Benchmarking Pocket-Scale Databases — *The TPC Technology Conference on Performance Evaluation & Benchmarking (TPC-TC)* 2019; Los Angeles, CA, USA; 16 pages, acceptance rate not available

W. Spoth, **T. Xie**, **O. Kennedy**, **Y. Yang**, B. Hammerschmidt, Z. H. Liu, D. Gawlick; SchemaDrill: Interactive Semi-Structured Schema Design — *Workshop on Human-In-the-Loop Data Analytics (HILDA)* 2018; Houston, TX, USA; pages, acceptance rate unknown; average rate for past conferences approximately 50.0%

P. Kumari, **O. Kennedy**; The Good and Bad Data — *North East Database Day (NEDB)* 2018; Boston, MA, USA; 2 pages, acceptance rate 19.64%

P. Kumari, S. Achmiz, **O. Kennedy**; Communicating Data Quality in On-Demand Curation — *International Workshop on Quality in DataBases (QDB)* 2016; New Dehli, India; 4 pages, acceptance rate not available

J. Freire, B. Glavic, **O. Kennedy**, H. Mueller (*Authors listed in alphabetical order*); The Exception That Improves The Rule — *Workshop on Human-In-the-Loop Data Analytics (HILDA)* 2016; San Fransisco, CA, USA; 6 pages, acceptance rate 50.0%

X. Niu, B. Arab, D. Gawlick, Z. H. Liu, V. Krishnaswamy, **O. Kennedy**, B. Glavic; Provenance-aware Versioned Dataworkspaces — *USENIX Workshop on Theory and Practice of Provenance (TaPP)* 2016; McLean, VA, USA; 6 pages, acceptance rate unknown; average rate for past conferences approximately 58.82%

G. Kul, **D. T. Luong**, **T. Xie**, **P. Coonan**, V. Chandola, **O. Kennedy**, S. Upadhyaya; Ettu: Analyzing Query Intents in Corporate Databases — *Workshop on Empirical Research Methods in Information Security (ERMIS)* 2016; Montreal, Canada; 4 pages, acceptance rate not available

O. Kennedy, J. A. Ajay, G. Challen, L. Ziarek; Pocket Data: The Need for TPC-MOBILE — *The TPC Technology Conference on Performance Evaluation & Benchmarking (TPC-TC)* 2015; Big Island, HI, USA; 16 pages, acceptance rate not available

G. Challen, J. A. Ajay, N. DiRienzo, **O. Kennedy**, A. Maiti, A. Nandugudi, S. Shantharam, J. Shi, G. P. Srinivasa, L. Ziarek (*Authors listed in alphabetical order*); maybe We Should Enable More Uncertain Mobile App Programming — *The ACM Workshop on Mobile Computing Systems and Applications (HotMobile)* 2015; Santa Fe, NM, USA; 6 pages, acceptance rate 28.75%

O. Kennedy, **Y. Yang**, J. Chomicki, R. Fehling, Z. H. Liu, D. Gawlick (*Authors listed UB first, then Oracle*; Detecting the Temporal Context of Queries — *Enabling Real-Time Business Intelligence (BIRTE)* 2014; Hangzhou, China; 17 pages, acceptance rate unknown; average rate for past conferences approximately 45.0%

S. Agarwal, **D. Bellinger**, **O. Kennedy**, **A. Upadhyay**, L. Ziarek (*Authors listed in alphabetical order*); Monadic Logs for Collaborative Web Applications — *The International Workshop on the Web and Databases (WebDB)* 2013; New York, NY, USA; 6 pages, acceptance rate 25.0%

A. Shieh, D. Williams, K. Walsh, **O. Kennedy**, P. Reynolds, E. G. Sirer, F. Schneider; Nexus: A New Operating System for Building Trustworthy Applications — *NSDI Poster Session (NSDI-Poster)* 2007; Cambridge, MA, USA; poster, acceptance rate not available

P. Reynolds, **O. Kennedy**, E. G. Sirer, F. Schneider; XMon-BGP: Securing BGP Using Using External Security Monitors — *NSDI Poster Session (NSDI-Poster)* 2007; Cambridge, MA, USA; poster, acceptance rate not available

Book Chapters (1)

O. Kennedy, B. Glavic; Analyzing Uncertain Tabular Data — *Information Quality in Information Fusion and Decision Making* — Information Fusion and Data Science *2019* (34 pages, acceptance rate not available)

Technical Reports (10)

S. Feng, B. Glavic, **A. Huber**, **O. Kennedy**, A. Rudra (*Authors listed in alphabetical order*); Computing expected multiplicities for bag-TIDBs with bounded multiplicities — *The ArXiv (ArXiv),* 9 pages

D. Balakrishnan, L. Ziarek, **O. Kennedy**; Just-in-Time Index Compilation — *The ArXiv (ArXiv),* 24 pages

S. Feng, **A. Huber**, B. Glavic, **O. Kennedy**; Uncertainty Annotated Databases - A Lightweight Approach for Dealing with Uncertainty — *Illinois Inst. Tech. Technical Report (IIT-TR),* 20 pages

N. Meneghetti, **O. Kennedy**, W. Gatterbauer; Beta Probabilistic Databases: A Scalable Approach to Belief Updating and Parameter Learning — *SUNY Buffalo Computer Science & Engineering Report (UB CSE),* 12 pages

G. Kul, **D. T. Luong**, **T. Xie**, **P. Coonan**, V. Chandola, **O. Kennedy**, S. Upadhyaya; Summarizing Large Query Logs in Ettu — *The ArXiv (ArXiv),* 12 pages

A. Nandi, **Y. Yang**, **O. Kennedy**, B. Glavic, R. Fehling, Z. H. Liu, D. Gawlick; Mimir: Bringing CTables into Practice — *The ArXiv (ArXiv),* 12 pages

O. Kennedy, L. Ziarek; Bar_{QL}: Collaborating Through Change — *SUNY Buffalo Computer Science & Engineering Report (UB CSE) #2013-03,* 10 pages

O. Kennedy, L. Ziarek; Bar_{QL}: Collaborating through Change — *The ArXiv (ArXiv),* 9 pages

J. Yang, E. Vee, S. Vassilvitskii, J. Tomlin, J. Shanmugasundaram, **O. Kennedy**; Inventory Allocation for Online Graphical Display Advertising — *The ArXiv (ArXiv),* 25 pages

O. Kennedy, C. Koch, A. Demers; Dynamic approaches to in-network aggregation — *The ArXiv (ArXiv),* 11 pages

Patents (2)

S. Nath, S. Lee, S. Smyl, C. Loboz, **O. Kennedy**; Efficient Optimization over Uncertain Data — USA Patent #US8650180B2

O. Kennedy, T. Anastasakos, J. Mao, E. Vee, J. Shanmugasundaram, S. Vassilvitskii, J. Yang, J. Tomlin; System for Display Advertising Optimization Using Click or Conversion Performance — USA Patent #US8311885B2

Artifacts (5)

Vizier (software) First released January 2018
An interactive data-centric notebook.
`https://vizierdb.info`

Just-in-Time Data Structures (software) First released July 2014
An index structure that dynamically adapts to variable workloads.
`http://github.com/UBOdin/jitd`
As of July 27, 2017, Just-in-Time Data Structures had 5 GitHub watchers, 5 GitHub stars, and 16 GitHub forks

Mimir (software) First released January 2015
A probabilistic overlay for database systems
`http://mimirdb.info`
As of July 27, 2017, Mimir had 11 GitHub watchers, 8 GitHub stars, and 8 GitHub forks

DBToaster (software) First released May 2012
A tool that dynamically compiles database engines optimized for any given query workload.
`http://www.dbtoaster.org`
As of December 10, 2014, DBToaster had 4191 unique website visitors and 1032 unique downloads

PigOut (software) First released October 2014
A multi-cluster runtime for Pig Latin
`http://pigout.cse.buffalo.edu`