

# Matthew Butrovich

Last updated: March, 2023

Email: [mbutrovi@cs.cmu.edu](mailto:mbutrovi@cs.cmu.edu)

Website: [mattbutrovi.ch](http://mattbutrovi.ch)

5000 Forbes Avenue

Pittsburgh, PA 15213

## Education



- 2019 – on      **Ph.D., Computer Science**, Carnegie Mellon University
- 2017 – 2018    **M.Sc., Computer Science**, Carnegie Mellon University
- 2013 – 2017    **B.Sc., cum laude, Computer Science and Engineering**, University of California, Irvine
- 2010 – 2013    **Certificate of Achievement, Computer Science**, Irvine Valley College

## Publications

### Peer-reviewed Papers

- 2023      Wan Shen Lim, **Matthew Butrovich**, William Zhang, Andrew Crotty, Lin Ma, Peijing Xu, Johannes Gehrke, and Andrew Pavlo. “Database Gyms”. *CIDR*. 
- 2022      **Matthew Butrovich**, Wan Shen Lim, Lin Ma, John Rollinson, William Zhang, Yu Xia, and Andrew Pavlo. “Tastes Great! Less Filling! High Performance and Accurate Training Data Collection for Self-Driving Database Management Systems”. *SIGMOD Conference*. 
- 2022      Yu Xia, Xiangyao Yu, **Matthew Butrovich**, Andrew Pavlo, and Srinivas Devadas. “Litmus: Towards a Practical Database Management System with Verifiable ACID Properties and Transaction Correctness”. *SIGMOD Conference*. 
- 2021      Lin Ma, William Zhang, Jie Jiao, Wuwen Wang, **Matthew Butrovich**, Wan Shen Lim, Prashanth Menon, and Andrew Pavlo. “MB2: Decomposed Behavior Modeling for Self-Driving Database Management Systems”. *SIGMOD Conference*. 
- 2021      Amadou Ngom, Prashanth Menon, **Matthew Butrovich**, Lin Ma, Wan Shen Lim, Todd C. Mowry, and Andrew Pavlo. “Filter Representation in Vectorized Query Execution”. *DaMoN*. 
- 2021      Ling Zhang, **Matthew Butrovich**, Tianyu Li, Andrew Pavlo, Yash Nannapaneni, John Rollinson, Huanchen Zhang, Ambarish Balakumar, Daniel Biales, Ziqi Dong, Emmanuel J. Eppinger, Jordi E. Gonzalez, Wan Shen Lim, Jianqiao Liu, Lin Ma, Prashanth Menon, Soumil Mukherjee, Tanuj Nayak, Amadou Ngom, Dong Niu, Deepayan Patra, Poojita Raj, Stephanie Wang, Wuwen Wang, Yao Yu, and William Zhang. “Everything is a Transaction: Unifying Logical Concurrency Control and Physical Data Structure Maintenance in Database Management Systems”. *CIDR*. 
- 2020      Tianyu Li, **Matthew Butrovich**, Amadou Ngom, Wan Shen Lim, Wes McKinney, and Andrew Pavlo. “Mainlining Databases: Supporting Fast Transactional Workloads on Universal Columnar Data File Formats”. *Proc. VLDB Endow.* 

### Miscellaneous Writings

- 2021      Andrew Pavlo, **Matthew Butrovich**, Lin Ma, Prashanth Menon, Wan Shen Lim, Dana Van Aken, and William Zhang. “Make Your Database System Dream of Electric Sheep: Towards Self-Driving Operation”. *Proc. VLDB Endow.* 
- 2019      Andrew Pavlo, **Matthew Butrovich**, Ananya Joshi, Lin Ma, Prashanth Menon, Dana Van Aken, Lisa Lee, and Ruslan Salakhutdinov. “External vs. Internal: An Essay on Machine Learning Agents for Autonomous Database Management Systems”. *IEEE Data Eng. Bull.* 

## Professional Experience

---

2009 – 2017    Technician V, *Western Digital Corporation*.

## Teaching Experience

---

2022	15-799: Special Topics in Database Systems Head teaching assistant <i>Carnegie Mellon University</i>	<a href="#">🔗</a>
2021	15-445/645: Database Systems Teaching assistant <i>Carnegie Mellon University</i>	<a href="#">🔗</a>
2020	15-721: Advanced Database Systems Head teaching assistant <i>Carnegie Mellon University</i>	<a href="#">🔗</a>
2019	15-445/645: Database Systems Lecture on two-phase locking concurrency control <i>Carnegie Mellon University</i>	<a href="#">🔗</a> <a href="#">▶</a>
2016	CSE145A/145B: Embedded Software and Laboratory Teaching assistant <i>University of California, Irvine</i>	
2015	CSE46: Data Structure Implementation and Analysis Lab tutor <i>University of California, Irvine</i>	
2014 – 2015	CSE45C: Programming in C/C++ as a Second Language Lab tutor <i>University of California, Irvine</i>	

## Academic Service

---

### To the University

2022 – on	Diversity, Equity and Inclusion Committee, Computer Science Department, <i>Carnegie Mellon University</i> .
2020 – on	Doctoral Review Committee, Computer Science Department, <i>Carnegie Mellon University</i> .
2019 – on	MSCS Admissions Committee, Computer Science Department, <i>Carnegie Mellon University</i> .






### To the Profession

2023	Reviewer, <i>ACM SIGMOD Artifacts &amp; Reproducibility Committee</i> .	<a href="#">🔗</a>
2022	Organizer, <i>¡Databases! - A Database Seminar Series</i> .	<a href="#">🔗</a>
2022	Organizer, <i>Vaccination Database Tech Talks - Booster</i> .	<a href="#">🔗</a>
2021	Organizer, <i>Vaccination Database Tech Talks - Second Dose</i> .	<a href="#">🔗</a>
2021	Organizer, <i>Vaccination Database Tech Talks - First Dose</i> .	<a href="#">🔗</a>
2020	Organizer, <i>Quarantine Database Tech Talks</i> .	<a href="#">🔗</a>


## Presentations

---

### Talks





- 2023 **Matthew Butrovich** and Andrew Pavlo. Embedding Database Logic in the Operating System Is Finally a Good Idea, *NEDB Day*, Boston, MA. 
- 2022 **Matthew Butrovich**, Karthik Ramanan, John Rollinson, Wan Shen Lim, William Zhang, Justine Sherry, and Andrew Pavlo. Tigger: A Database Proxy That Bounces With User-Bypass, *PDL Retreat*, Pittsburgh, PA. 
- 2022 **Matthew Butrovich** and Andrew Pavlo. The Sugar-free Chocolate of Databases, *HPTS*, Asilomar, CA. 
- 2022 **Matthew Butrovich** and Andrew Pavlo. Tastes Great! Less Filling! High Performance and Accurate Training Data Collection for Self-Driving Database Management Systems, *SIGMOD*, Philadelphia, PA.  

### Posters

- 2023 **Matthew Butrovich**, Wan Shen Lim, Lin Ma, John Rollinson, William Zhang, Yu Xia, and Andrew Pavlo. Tastes Great! Less Filling! High Performance and Accurate Training Data Collection for Self-Driving Database Management Systems, *NEDB Day*, Boston, MA. 
- 2022 **Matthew Butrovich**, Wan Shen Lim, Lin Ma, John Rollinson, William Zhang, Yu Xia, and Andrew Pavlo. High Performance and Accurate Training Data Collection for Self-Driving Database Management Systems, *PDL Retreat*, Pittsburgh, PA.
- 2022 Lin Ma, William Zhang, Jie Jiao, Wuwen Wang, **Matthew Butrovich**, Wan Shen Lim, Prashanth Menon, and Andrew Pavlo. MB2: Decomposed Behavior Modeling for Self-Driving Database Management Systems, *PDL Retreat*, Pittsburgh, PA.
- 2020 **Matthew Butrovich**, ... and Andrew Pavlo. NoisePage: The Self-Driving Database Management System, *PDL Visit Day*, Pittsburgh, PA.
- 2019 **Matthew Butrovich** and Andrew Pavlo. Self-Driving DBMS Engineering, *HPTS*, Asilomar, CA.
- 2019 **Matthew Butrovich** and Andrew Pavlo. Self-Driving DBMS Engineering, *PDL Visit Day*, Pittsburgh, PA.
- 2018 Tianyu Li, **Matthew Butrovich**, and Andrew Pavlo. Mainlining Databases: Supporting Fast Transactional Workloads on Universal Columnar Data File Formats, *PDL Retreat*, Bedford Springs, PA.

## Open-source Software

---

- 2018 – 2022 **NoisePage**    
 PostgreSQL-compatible in-memory DBMS designed for autonomous operation.  
 Role: Founding developer responsible for HyPer-style MVCC, PostgreSQL protocol, training data collection, and project management.
- 2017 – 2018 **Peloton**    
 PostgreSQL-compatible in-memory DBMS designed for autonomous operation.  
 Role: Developer optimizing concurrency control and garbage collection.