

# Ryan Marcus

---

Ph.D. Candidate  
Brandeis University

Phone: 505-500-4365  
Email: [ryan@cs.brandeis.edu](mailto:ryan@cs.brandeis.edu)  
Web: <https://rmarcus.info>  
Twitter: @RyanMarcus  
GitHub: RyanMarcus

## RESEARCH INTERESTS

Cloud databases, machine learning, parallel and distributed programming, high performance computing, programming languages & compilers.

## EDUCATION

|  |                  |
|--|------------------|
| <b>Brandeis University Ph.D. Candidate, Computer Science</b>   | 2016 – Current   |
| Advisor: Olga Papaemmanouil (expected graduation: Spring 2019) | Waltham, MA. USA |
| <b>Brandeis University M.A. Computer Science</b>               | January 2016     |
| <b>University of Arizona</b>                                   | May 2014         |

## EMPLOYMENT

|                           |                  |
|---------------------------|------------------|
| <b>Research assistant</b> | 2015 – Current   |
| Brandeis University       | Waltham, MA. USA |

WiSeDB: a learning-based cost and performance management service for cloud databases. Supervised and reinforcement learning techniques to automate resource provisioning and workload scheduling by generating models that minimize monetary cost, while supporting performance SLAs.

|                     |                 |
|---------------------|-----------------|
| <b>Ph.D. Intern</b> | Summer 2018     |
| Google              | Irvine, CA. USA |

Improving ad spend and revenue modeling for the Attribution 360 team. Developed convergence detection method to speed up model training 33%. Developed aggregation methods to model performance in low-variance scenarios.

|                        |                   |
|------------------------|-------------------|
| <b>Research intern</b> | Summer 2017       |
| Microsoft Research     | Richmond, WA. USA |

Machine learning techniques to predict query performance regressions due to physical design changes. Online transfer learning via learned kernels to adapt models to new datasets.

|  |                    |
|--|--------------------|
| <b>Software developer</b>                | Summer 2016        |
| Hewlett-Packard Enterprise (HPE) Vertica | Cambridge, MA. USA |

Design and implementation of improved query plan execution algorithms, achieving a 4x performance boost for rollup analytic queries.

## High performance computing (HPC) engineer

Los Alamos National Laboratory

2009-2015

Los Alamos, NM, USA

Developed a machine-learning framework for automatic performance analysis, currently used to analyze high-performance hydrodynamics codes. Parallelized serial algorithms on GPUs. Designed novel median filter approximation algorithm, improved 3D image reconstruction performance by 3x. Implemented exascale co-design neutron transport code.

## HONORS & AWARDS

Computer Science Outstanding Teaching Fellow Award

Brandeis University, 2015

Michtom Fellowship

Brandeis University, 2014

CEDA Summa Cum Laude National Debate Scholar

Idaho State University, 2013

## PUBLICATIONS

1. Ryan Marcus, Olga Papaemmanouil. "Towards a Hands-Free Query Optimizer through Deep Learning." 9th Biennial Conference on Innovative Data Systems Research (**CIDR 19**) 2019.
2. Ryan Marcus Olga Papaemmanouil. "Deep Reinforcement Learning for Join Order Enumeration." First International Workshop on Exploiting Artificial Intelligence Techniques for Data Management (**aiDM@SIGMOD 18**) 2018.
3. Ryan Marcus Olga Papaemmanouil Sofiya Semenova Solomon Garber. "NashDB: An Economic Approach to Fragmentation, Replication and Provisioning for Elastic Databases." 37th ACM Special Interest Group in Data Management (**SIGMOD 18**) 2018.
4. Ryan Marcus, Sofiya Semenova, and Olga Papaemmanouil. "A Learning-based Service for Cost and Performance Management of Cloud Databases (Demo)." IEEE International Conference on Data Engineering (**ICDE 17**) 2017.
5. Ryan Marcus and Olga Papaemmanouil. "Releasing Cloud Databases from the Chains of Performance Prediction Models." 8<sup>th</sup> Biennial Conference on Innovative Data Systems Research (**CIDR 17**) 2017.
6. Ryan Marcus and Olga Papaemmanouil. "WiSeDB: a learning-based workload management advisor for cloud databases." VLDB Endowment (**VLDB 16**), Volume 9, Issue 10, June 2016.
7. Ryan Marcus and Olga Papaemmanouil. "Workload Management for Cloud Databases via Machine Learning." Workshop on Cloud Data Management and the IEEE International Conference on Data Engineering (**CloudDM@ICDE 16**), 2016.

## NON-REFEREED PUBLICATIONS

1. Ryan Marcus, William Nystrom, David Gunter and Cornell Wright. "Shared Memory for Many-Core Hydrodynamics Codes", Tech. Rep. LA-UR-15-25303, Los Alamos Nat. Lab., 2015.
2. Ryan Marcus. "Techniques for Automated Performance Analysis." Tech. Rep. LA-UR-14-26577, Los Alamos Nat. Lab., 2014.
3. Ryan Marcus and William Ward. "DP: A fast median filter approximation." Tech. Rep. LA-UR-13-25331, Los Alamos Nat. Lab., 2013.

4. Ryan Marcus. "MCMini: Monte Carlo on GPGPU." Tech. Rep. LA-UR-12-23206, Los Alamos Nat. Lab., 2012.
5. Lawrence Cox and Ryan Marcus. "Developing a Monte Carlo mini-app for exascale co-design." Tech. Rep. LA-UR-11-06085, Los Alamos Nat. Lab., 2011 (industry demo at SC11).

## TEACHING EXPERIENCE

|  |                         |
|--|-------------------------|
| T.A. for Networked Information Systems   | Spring 2018             |
| T.A. for Advanced Programming Techniques | Spring 2016 – Fall 2017 |
| T.A. for Databases Management Systems    | Fall 2015               |
| T.A. for Theory of Computation           | Spring 2015             |
| T.A. for Operating Systems               | Fall 2014               |

## INVITED TALKS

|   |                             |
|---|-----------------------------|
| <i>ReJOIN: Deep Reinforcement Learning for Join Order Enumeration</i><br>IBM Research Lab @ 1 Rogers Street         | 2018<br>Cambridge, MA. USA  |
| <i>Deep Learning Applications to Database Query Optimization</i><br>University of Wisconsin-Madison                 | 2018<br>Madison, WI. USA    |
| <i>Deep Learning Meets Query Optimization</i><br>Massachusetts Institute of Technology                              | 2018<br>Cambridge, MA. USA  |
| <i>Machine Learning Powered Clouds</i><br>Harvard University  | 2018<br>Cambridge, MA. USA  |
| <i>Machine Learning for Cloud Database Workload Management</i><br>Hewlett-Packard Enterprise (HPE) Vertica          | 2016<br>Cambridge, MA. USA  |
| <i>Machine Learning for Humans</i><br>Los Alamos National Laboratory Data Science Institute                         | 2015<br>Los Alamos, NM. USA |
| <i>Shared Memory for Many-Core: A Hydrodynamics Case Study</i><br>Los Alamos National Laboratory Research Symposium | 2015<br>Los Alamos, NM. USA |

## SKILLS

Java, Python 2/3, C/C++, JavaScript (ES5/6), Fortran, Weka, MapReduce, Hadoop, Keras, scikit-learn, PyTorch, VueJS, NodeJS, Express, Emscripten, PostgreSQL, SQLite, MariaDB/MySQL, Vertica.

## PROFESSIONAL MEMBERSHIP

- IEEE Student Member
- ACM Student Member

## PERSONAL

I have gone to a few hackathons and created a few personal programming projects. You can see them all on my website: <https://rmarcus.info/blog/pubs/#projects>

I also organize the Brandeis Computer Science Ph.D. seminars: <http://rm.cab/seminars>