

# DR RENATA BOROVIKA-GAJIC

CV

Senior Lecturer in Data Analytics  
The University of Melbourne  
School of Computing and Information Systems

**Email:** renata.borovica@unimelb.edu.au  
**Website:** www.renata.borovica-gajic.com  
**Linkedin:** http://ch.linkedin.com/in/renataborovica

## EDUCATION

<b>2018 - 2019</b>	The University of Melbourne, Melbourne, Australia <b>Graduate Certificate in University Teaching (GCUT)</b>
<b>2010 – 2016</b>	Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland <b>PhD in Computer Science</b> Thesis: Toward timely, predictable and cost-effective data analytics
<b>2003 – 2008</b>	Faculty of Technical Sciences, University of Novi Sad, Serbia <b>Master in Electrical and Computer Engineering</b> (integrated MSc and BSc studies) Thesis: Historical Server Query Performance Improvement Based on Grid Processing of Dynamic Data <b>National award for the best MSc thesis in the field of mathematics and computer sciences</b>

## PRESENT AND PAST APPOINTMENTS

<b>2017-now</b>	<b>The University of Melbourne (UoM), School of Computing and Information Systems (Australia)</b> Senior Lecturer in Data Analytics
<b>2010-2016</b>	<b>Swiss Federal Institute of Technology EPFL (Lausanne, Switzerland)</b> Research assistant in the Data-Intensive Applications and Systems Laboratory (DIAS) working under the supervision of Professor Anastasia Ailamaki
<b>2013-2013</b>	<b>MICROSOFT Research (Redmond, WA, USA), duration 3 months</b> Research Intern in the Data Management, Exploration and Mining Group supervised by Dr. Surajit Chaudhuri, Dr. Vivek Narasayya and Dr. Christian Konig
<b>2005-2010</b>	<b>DMS Group LTD (Novi Sad, Serbia)</b> Senior Software Engineer (Database Team)
<b>2008-2009</b>	<b>SIEMENS Energy Inc. (Minneapolis, Minnesota, USA)</b> Software Developer, Consultant (Database Team)

## AWARDS

- Excellence in Teaching and Learning for Early Career Academics, Faculty of Engineering and IT (FEIT), UoM, 2020.
- Excellence in Teaching and Learning, School of Computing and Information Systems (CIS), UoM, 2018.
- School of Computer and Communication Sciences achievement award, EPFL, 2015 (CHF 1,000).
- 2015 Travel award for attending the International Conference on Data Engineering (ICDE) (USD 1,000).
- Dositeja Fellowship from the Serbian Government for students studying abroad, 2012-2013 (EU 7,000).
- School of Computer and Communication Sciences, fellowship, EPFL, Switzerland (CHF 55,500)
- Serbian national award 'Mileva Maric Einstein' for the best master thesis in the field of mathematics and computer sciences in 2008.
- Best graduation student at the Faculty of Technical Sciences in 2008.
- Recognized as one of the best 100 students in Serbia by the Government of Serbia in 2007 (EU 1,000).
- DMS Group Ltd. fellowship student (mentoring program), 2005-2008 (three years salary).

## FUNDING

- Google Award for Inclusion Research, 2021. (USD 60,000)
- FEIT Interdisciplinary Grant, The University of Melbourne, 2021. (AU 30,000)
- Australian Research Data Commons (ARDC), Platforms, 2020. (AUD 925,000)
- Melbourne Networked Society Institute, Seed funding, 2018. (AUD 28,000)
- Early Career Researcher (ECR) Grant, The University of Melbourne, 2018. (AUD 39,000)

## RESEARCH RECOGNITION

---

- **Academic:** My work has repeatedly been published at three premier database conferences (SIGMOD, VLDB, ICDE, all CORE A\*) and has been internationally recognised as pioneering. Five years post-PhD (with significant academic interruptions), I have published *25 peer-reviewed publications*, and my work has attracted 411 citations to date. My publications include 12 full conference proceedings papers, 8 journal publications, one book chapter, 2 workshops papers, and 1 patent. My work has been cited by faculty members at top computer science schools such as Harvard, Yale, University of Maryland, and highlighted during conference keynotes at the top three database conferences (e.g., ICDE 2014, SIGMOD 2014, VLDB 2015, VLDB 2017), further demonstrating international recognition. Furthermore, my work has been invited to the Research Highlights Section of the Communication of ACM (impact factor 6.469), a recognition awarded to a small number of cutting-edge research papers in computer science. For the past 5 years, my performance indicators are substantially higher than the average researcher from my field: my field-weighted citation impact is 3.07, which is 207% higher than the expected average for researchers in my field; 50% of my publications are in the top 10% journals ranked world-wide; and 11.1% are in the top 1% most cited publications world-wide (from Scopus, 2021). I am on a fast upward trajectory, with 11 publications accepted in 2020/2021, 8 of those ranked A/A\*.
- **Industrial:** The concepts and ideas proposed in my work have been adopted by the leading database vendors such as Microsoft, showcasing strong industrial impact. Furthermore, my work has been patented and underpinned a successful start-up located in Switzerland.

## RESEARCH OUTPUTS (ORCID ID: 0000-0003-3503-4123, <https://scholar.google.ch/citations?user=HYkhTeYAAAAJ>)

---

1. R. Muralidhar, R. Borovica-Gajic, R. Buyya: Energy Efficient Computing Systems: Architectures, Abstractions & Modeling to Techniques and Standards. ACM Computing Surveys (**CSUR**), full length article, p. 1-63, to appear, 2022. **(Core A)**
2. B. Oetomo, M. Perera, R. Borovica-Gajic, B.I.P. Rubinstein: Cutting to the Chase with Warm-Start Contextual Bandits. The IEEE International Conference on Data Mining (**ICDM**), full length published paper, p. to appear, 2021. **(Core A\*)**
3. R. Chittor Sundaram, E. Naghizade, R. Borovica-Gajic, M. Tomko: Can you fixme? An intrinsic classification of contributor-identified spatial data issues using topic models. International Journal of Geographical Information Science (**IJGIS**), full length published article, p. 1-30, 2021. **(Core A)**
4. M. Perera, B. Oetomo, B.I.P. Rubinstein, R. Borovica-Gajic: DBA bandits: Self-driving index tuning under ad-hoc, analytical workloads with safety guarantees. The 37th IEEE International Conference on Data Engineering (**ICDE**), full length published paper, p. 600-611, 2021. **(Core A\*)**
5. R. Chittor Sundaram, E. Naghizade, R. Borovica-Gajic, M. Tomko: Harnessing spatio-temporal patterns in data for nominal attribute imputation. Transactions in GIS (**TGIS**), full length published article, vol. 24(4), p. 1001-1032, 2020. **(Core B)**
6. Y. Xu, L. Kulik, R. Borovica-Gajic, A. AlDwyish, J. Qi: Highly Efficient and Scalable Multi-hop Ride-sharing. The 28th International Conference on Advances in Geographic Information Systems (**SIGSPATIAL**), full length published paper, p. 215-226, 2020. **(Core A)**
7. M. Li, F. M. Choudhury, R. Borovica-Gajic, Z. Wang, J. Xin, J. Li: CrashSim: An Efficient Algorithm for Computing SimRank over Static and Temporal Graphs. The 36th IEEE International Conference on Data Engineering (**ICDE**), full length published paper, p. 1141-1152, 2020. **(Core A\*)**
8. U. Gunarathna, H. Xie, E. Tanin, S. Karunasekera, R. Borovica-Gajic: Real-Time Lane Configuration with Coordinated Reinforcement Learning. Machine Learning and Knowledge Discovery in Databases: Applied Data Science Track - European Conference (**ECML/PKDD**), full length published paper, vol. (4), p. 291-307, 2020. **(Core A)**
9. Y. Xu, J. Qi, Renata Borovica-Gajic, Lars Kulik: GeoPrune: Efficiently Matching Trips in Ride-sharing Through Geometric Properties. The 32nd International Conference on Scientific and Statistical Database Management (**SSDBM**), full length published paper, vol. (5), p. 1-12, 2020. **(Core A)**
10. R. Borovica-Gajic, J. Qi, W. Wang: Databases Theory and Applications - 31st Australasian Database Conference, ADC 2020, Proceedings. Lecture Notes in Computer Science 12008, 2020.
11. N. F. Setiawan, B. I. P. Rubinstein, and R. Borovica-Gajic: Function Interpolation for Learned Index Structures. Australian Database Conference (**ADC**), full length published paper, p. 68-80, 2020. **(best paper honorable mention)**

12. **(Book chapter)** R. Borovica-Gajic and I. Alagiannis, *Database System Concepts*, 7th edition, by Silberschatz, Korth and Sudarshan. Chapter 32 on PostgreSQL, p. 1-58, 2019.
13. R. Appuswamy, G. Graefe, R. Borovica-Gajic, A. Ailamaki: The five-minute rule 30 years later and its impact on the storage hierarchy. **Communications of the ACM**, full length published article, vol. 62(11), p. 114-120, 2019. **(Core A)**
14. R. Borovica-Gajic, S. Idreos, A. Ailamaki, M. Zukowski and C. Fraser: Smooth Scan: Robust Access Path Selection without Cardinality Estimation, *The VLDB Journal*, full length published article, vol. 27(4), p. 521-545, 2018. **(Core A\*)**
15. Y. Xu, J. Qi, R. Borovica-Gajic, L. Kulik: Finding All Nearest Neighbors with a Single Graph Traversal. Database Systems for Advanced Applications - 23rd International Conference (**DASFAA**), full length published article, p. 221-238, 2018. (Core B)
16. R. Borovica-Gajic, G. Graefe, A. Lee: Robust Performance in Database Query Processing (Dagstuhl Seminar 17222). **Dagstuhl Reports** 7(5), p. 169-180, 2017.
17. R. Appuswamy, R. Borovica-Gajic, G. Graefe, A. Ailamaki: The Five-minute Rule Thirty Years Later and its Impact on the Storage Hierarchy. The 8th International Workshop on Accelerating Analytics and Data Management Systems Using Modern Processor and Storage Architectures (**ADMS@VLDB**), full length published paper, p. 1-8, 2017.
18. R. Borovica-Gajic: Toward Timely, Predictable and Cost-effective Data Analytics. PhD Thesis, 227 pages, EPFL, 2016. *Nominated for Roger Needham PhD Award in 2017 for an exceptional, innovative contribution to knowledge in the systems area.*
19. A. Ailamaki, S. Idreos, I. Alagiannis, R. Borovica, and M. Branco: Query Management System and Engine Allowing for Efficient Query Execution on Raw Data Files. **US Patent** (US 9298754), 2016.
20. R. Borovica-Gajic, R. Appuswamy, and A. Ailamaki: Cheap Data Analytics Using Cold Storage Devices. The 42nd International Conference on Very Large Data Bases (**VLDB**), full length published paper, vol. 9(12), p. 1029-1040, 2016. **(Core A\*)**
21. R. Borovica-Gajic, S. Idreos, A. Ailamaki, M. Zukowski and C. Fraser: Smooth Scan: Statistics-oblivious Access Paths. The 31st International Conference on Data Engineering (**ICDE**), full length published paper, p.315-326, 2015. **(Core A\*)**
22. I. Alagiannis, R. Borovica-Gajic, M. Branco, S. Idreos, and A. Ailamaki: NoDB: Efficient Query Execution on Raw Data Files. **Communications of the ACM, Research Highlights**, full length published article, vol. 58(12), p. 112-121, 2015. **(Core A)**
23. I. Alagiannis, R. Borovica, M. Branco, S. Idreos and A. Ailamaki: NoDB in Action: Adaptive Query Processing on Raw Data (demo). The 38th International Conference on Very Large Data Bases (**VLDB**), full length published demonstration, vol. 5(12), p. 1942-1945, 2012. **(Core A\*)**
24. R. Borovica, I. Alagiannis and A. Ailamaki: Automated Physical Designers: What You See is (Not) What You Get. The 5th International Workshop on Testing Database Systems (**DBTest@SIGMOD**), full length published paper, 9 pages, 2012.
25. I. Alagiannis, R. Borovica, M. Branco, S. Idreos and A. Ailamaki: NoDB: Efficient Query Execution on Raw Data Files. ACM International Conference on Management of Data (**SIGMOD**), full length published paper, p. 241-252, 2012. **(Core A\*)**
26. T. Heinis, M. Branco, I. Alagiannis, R. Borovica, F. Tauheed and A. Ailamaki: Challenges and Opportunities in Self-Managing Scientific Databases. IEEE Data Engineering Bulletin (**DEB**), full length published article, vol. 34(4), p. 44-52, 2011.

## TEACHING EXPERIENCE

---

- **INFO20003 Database Systems (420 students) with Dr Ida Someh, The University of Melbourne, 2017 s2.**  
Revamped the subject adopting new syllabus and new lecture slides. I co-taught the subject with Dr Someh.
- **INFO20003 Database Systems (220 students), The University of Melbourne, 2018 s1.**  
Took over the subject teaching completely and developed new tutorial material and labs, inspired by industrial use cases. The tutorials adopted problem-based learning with a lot of exercises, while the labs used an industrial strength database system. Designed new assignments that were aligned with the new syllabus and subject learning objectives.
- **INFO20003 Database Systems (355 students), The University of Melbourne, 2018 s2.**  
Developed practice material with 5 additional case studies, and exam booklet to help students with the exam preparation. Introduced a Peerwise competition for additional exam preparation.
- **INFO20003 Database Systems (310 students), The University of Melbourne, 2019 s1.**  
Developed weekly quizzes and introduced marks feedback generator that sends personal feedback and assignment results (with a detailed breakdown of individual marks) to students over email.
- **INFO20003 Database Systems (478 students), The University of Melbourne, 2019 s2.**  
Used blended learning, and developed short polls, case studies and quiz questions all incorporated into the lectures. Developed a database playground module with visual representations of internal database mechanisms targeted for students of non-computer science background.
- **INFO20003 Database Systems (589 students), The University of Melbourne, 2020 s2.**  
Adopted flipped classroom model. Introduced mini-video sequences with quiz questions incorporated within. Mindful of the budget restrictions, developed an automated marking platform for the assignments that auto marked 2 out of 3 assignments that we have in this subject (reducing marking hours by approximately 50%).

## SUPERVISION

---

### PhD Students

1. Farzaneh Zirak: "Using Eye-Tracking Technology to Improve Visual Data Exploratory Tools Performance", PhD underway, co-supervision with Dr Farhana Choudhury. The University of Melbourne, 2022-.
2. Xinling Shen: "Dynamic Indoor Evacuation Planning in an Emergency", PhD underway, co-supervision with Prof Stephan Winter. The University of Melbourne, 2022-.
3. Renlord Yang: "An empirical analysis of the systemic performance of decentralized cryptocurrency systems", PhD underway, co-supervision with Prof Udaya Parampalli and A/Prof Toby Murray. The University of Melbourne, 2020-.
4. Rajeev Muralidhar: "Towards energy efficient systems", PhD underway, co-supervision with Prof Raj Buyya, The University of Melbourne, 2019-.
5. Udesh Gunarathna: "Dynamic network configuration based on user traffic patterns", PhD underway, co-supervision with Prof Egemen Tanin and Prof Shanika Karunasekera. The University of Melbourne, 2018-.
6. Malinga Perera: "Learned Approach for Self-Tuning Physical Design Structures Using Multi-Armed Bandit Algorithms", PhD underway, co-supervision with Prof Ben Rubinstein. The University of Melbourne, 2018-.
7. Bastian Oetomo: "Self-driving Databases with Multi-Armed Bandit Algorithms", PhD underway, co-supervision with Prof Ben Rubinstein. The University of Melbourne, 2018-.
8. Rajesh Chittor Sundaram: "Automated Cleaning of Spatial Data", PhD underway, co-supervision with Dr Martin Tomko and Dr Elham Naghi Zadeh. The University of Melbourne, 2017-.
9. Yixin Xu: "All Nearest Neighbor Applications", PhD *completed*, co-supervision with Prof Lars Kulik and Dr Jianzhong Qi. The University of Melbourne, 2017-2021.

### Master Students

- **The University of Melbourne:**  
Kai Zhang, Rida Fatima, Shashank Satish Parab, Yuqian Shi, Shreyash Patodia, Winston Angkajaya, Yun Chen, Arnab Rakshit, Zehua Zhang, Boheng Luan, Boyan Zhao, Shajid Mohammad, Jian Li, Changhui Qian, Naufal Fikri Setiawan, Peixin Lin, Zican Yao, Junyang Liu, Karan Bansal, Shenglan Yu, Kailun Huang, Wanmin Chen, Mengzhu Long, Xinze Li, Palak Somani, Linxuan Zhao, Sharda Poojita Pratti, Kumar Utkarsh, Sonny Theo Tumbur, Xinnan SHEN, Ye Yang, Mark Jones, Kazuya Soga, Minjian Chen, Tian Hang, Saransh Srivastava, Min Zhao, Zijie Pan, Aayush Mehta, Ze Liang, Abhijeet Singh
- **EPFL:**  
Alejandro Rivas, Dimitrios Sarigiannis, Lisa Zhou, Arnold Lacko, Alejandro Naser

## RECENT INVITED TALKS

---

- 2021, University of Iowa, Iowa, USA Data analytics for a penny
- 2021, University of Copenhagen, Copenhagen, Denmark Learning based algorithms within DBMS: Hype or Future
- 2020, Swiss Federal Institute of Technology, Lausanne, CH Multi-armed bandits stealing DBA jobs
- 2019, Australasian Database Conference, Sydney, AUS A tale of learning databases
- 2018, Data Science Colloquium, Brisbane, AUS Tips and tricks for effective PhD -- *Keynote*
- 2017, Dagstuhl, Germany Smooth Scan: One Access Path to Rule them All

## PROFESSIONAL SERVICE AND LEADERSHIP

---

### Leadership

- Diversity and Inclusion Lead, School of Computing and Information Systems, 2020-present
- SIGMOD RECORD Associate Editor (surveys), 2021-2023

### Conference and Seminar Organization

- 2022 ACM Conference on Management of Data (SIGMOD) CORE A\* Diversity and Inclusion Chair
- 2021 International Conference on Very Large Databases (VLDB) CORE A\* Registration Chair
- 2021 International Conference on Data Engineering (ICDE) CORE A\* PhD Symposium Chair
- 2021 International Symposium on Cloud Computing (SOCC) Registration Chair
- 2021 Dagstuhl Seminar on 'Database Indexing and Query Processing' Organizer
- 2020 Australasian Database Conference PC Chair
- 2017 Dagstuhl Seminar on 'Robust Performance in Database Query Processing' Organizer

### Recent Program Committee Member (*A\*/A denote flagship conferences*)

- 2022: SIGMOD (A\*), VLDB (A\*), ICDE (A\*), CIDR (A), PhDWorkshop@VLDB, AIDM@SIGMOD, DEEM@SIGMOD, DBTest@SIGMOD
- 2021: SIGMOD (A\*), EDBT 2021 (A)
- 2020: SIGMOD (A\*), VLDB (A\*), ICDE (A\*), DBTest@SIGMOD
- 2019: SIGMOD research & demo track (A\*), ICDE (A\*), VLDB (A\*), ADC, WomENCourage (women in STEM), CIKM (A), Information Systems (ERA A\*), Grace Hopper Celebration (women in STEM)