

Victor Giannakouris

- vg292@cornell.edu • (+1) 917 945 6484
- Google Scholar. Citations: 87, h-index: 4
- [LinkedIn](#)

- INTERESTS** I am interested in database systems, and especially in query optimization for federated query execution engines. I have research and industry experience on implementing database system components, ranging from indexes/storage engines (TileDB, MonetDB) to federated query optimizers for Spark SQL. Lately I have been working on leveraging machine learning in federated query optimization.
- EDUCATION**
- Cornell University**, Ithaca, NY, USA
- Ph.D., Computer Science 2021 – 2026
 - Area: Database Systems & Machine Learning
 - Advisor: Immanuel Trummer
- National and Kapodistrian University of Athens**, Athens, Greece
- M.Sc., Computing Systems Oct 2018 – Jul 2021
 - Thesis: An Online Sample Index for Approximate Query Processing
 - Advisor: Yannis Ioannidis
- Ionian University, Department of Informatics**, Corfu, Greece
- B.Sc., Informatics Sep 2012- May 2017
 - Thesis: Optimization of Distributed SQL Query Execution Over Multiple-Engine Environments
 - Advisor: Dimitrios Tsoumakos
- EMPLOYMENT**
- Cornell University**, Ithaca, NY, USA
- Graduate Research Assistant Aug 2021 – Present
 - TA: CS 4321 - Database Systems Implementation (Fall 2022)
 - TA: CS 4410 - Introduction to Operating Systems (Fall 2021, Spring 2022)
- IBM Research**, Almaden, CA, US
- Research Intern May 2022 – Aug 2022
 - Database Internals
- TileDB**, Boston, MA, US
- Senior Software Engineer Nov 2019 – Jul 2021
 - Database Internals, Storage Engines
- Unravel Data Systems**, Palo Alto, CA, US
- Software Engineer Oct 2018 – Oct 2019
 - Big Data Application Performance Management (APM)
 - Developed Unravel Platform's Cloud Migration component ([link](#))
- AISERA**, Palo Alto, CA, US
- Data Engineer Apr 2018 – Sep 2018
 - Development of data collection microservices on the company's platform's data ingestion component.
 - Tools: Kubernetes, AWS, Google Cloud Platform, Java, Python
- BEAT**, Athens
- Big Data Software Engineer Sep 2017 – Apr 2018
 - Development of streaming and ETL jobs over large-scale data.
 - Tuning and optimization of Spark SQL batch jobs on query and storage level.
 - Tools: Spark, YARN, Hadoop, Hive, Presto, Scala, Python
- Computing Systems Laboratory (CSLab)**, NTUA, Greece
- Researcher Jul 2015 – Aug 2017
 - R&D on Big Data Management, Databases and Multi-Engine Analytics.
 - Participated in four publications.
 - Development of a middleware optimization framework for Spark SQL, achieving execution time improvement for queries over heterogeneous environments up to one order of magnitude.
 - Participated in the EU Funded Project ASAP(A Scalable Analytics Platform).
- isMOOD Data Technology Services**
- Software Engineer Apr 2014 – Sep 2015
 - Data collection and integration into a unified source
 - Development, scheduling & optimization of web crawlers/spiders

