

Editor's Notes

Welcome to the September 2022 issue of the ACM SIGMOD Record!

This issue starts with the Database Principles column featuring an article by Arenas, Croquevielle, Jayaram, and Riveros. The article focuses on counting the answers to queries, a fundamental problem with applications in query evaluation, optimization, and visualization. In the face of the #P-hardness of the general case, the authors present recent results on approximate counting with fully polynomial-time randomized approximation schemes, and discuss applications of the results to automata theory, graph databases, and conjunctive-query evaluation.

The Surveys column presents an article by Azcoitia and Laoutaris. The article shares results and conclusions from a survey covering almost 200 commercial data-trading entities, with a discussion of the types of data that they trade and of their business models and the technologies that they rely upon. The authors also point to promising open research questions in the area, encouraging the community to contribute to benefiting the growing ecosystem of data-trading entities.

The Research column presents an article by Kennedy, Subramaniam, Galhotra, and Fernandez. In the domain of data markets, the authors focus on the problem of matching sellers with buyers for effective data allocation. The article presents the results of analyzing the data sold in the most popular data-marketplace platforms, and discusses the outcomes of interviewing both buyers and sellers toward understanding their experiences. The authors connect the outcomes of the buyer and seller interviews to the major challenges to allowing effective spread of data value in data marketplaces, and offer avenues of future research.

The Reminiscences on Influential Papers column features contributions by Fatma Özcan and Yuanyuan Tian.

The Advice to Mid-Career Researchers column features a contribution by Anastasia Ailamaki entitled “The Formidable Mid-Career Crisis.” The article shares advice to mid-career professionals on going after a risky vision, fostering collaborations, organizing and promoting their students, being picky with the service they accept, all while having a lot of fun in the process.

The Future of Data(base) Education column focuses on presenting a series of perspectives on data (base/science) education to help educators think about what we should be teaching in our courses, and what resources we should use to teach them. In this issue, the column features a contribution by Kumar, Deutsch, Gupta, Papakonstantinou, Salimi, and Vianu about the Database Lab at UC San Diego (UCSD). Three years ago, UCSD launched the Halicioglu Data Science Institute (HDSI) to shape the future of Data Science research, education and societal impact. As a result, the Database Lab in the Department of Computer Science and Engineering (CSE) took a holistic view of database education across the university to understand how database courses in CSE should interact with and/or be redesigned for use within HDSI as well as within the Rady School of Management. The article gives an overview of the resulting DB curricula, and their rationales.

The DBrainstorming column, whose goal is to discuss new and potentially controversial ideas that might be of interest and potentially of benefit to the research community, features in this issue an article by Balmau discussing the challenge of efficient data access from persistent storage in machine learning. The author argues for the need to characterize I/O patterns in machine learning with a focus on data preprocessing and training, and proposes to create the first open-access storage-focused

benchmark for machine learning. The article provides an overview of the approach and of the expected impact.

The Industry Perspectives column features an article by Jindal and Leeka that presents arguments in favor of considering query optimization as a service in modern cloud architectures. The authors introduce Oasis, a reference architecture for query optimizer as a service, and describe their success in developing an early version of Oasis. They also discuss the risks and responsibilities involved with Oasis to ensure it is a win-win for everyone.

The Reports column features an article by Kumar, Halevy, and Tatbul that discusses the new VLDB Scalable Data Science (SDS) research-track category. The goal of SDS is to attract cutting-edge and impactful real-world work in the scalable data-science arena to enhance the impact and visibility of the VLDB community on data-science practice, spur new technical connections, and inspire new follow-on research. In this report, the authors reflect on the inaugural year of SDS, provide relevant statistics, and share observations, lessons, and tips as inaugural Associate Editors for SDS.

The issue closes with a call for SIGMOD 2024 papers, with an announcement on the planned quarterly deadlines for the research track of the conference.

On behalf of the SIGMOD Record Editorial board, I hope that you enjoy reading the September 2022 issue of the SIGMOD Record!

Your submissions to the SIGMOD Record are welcome via the submission site:

<https://mc.manuscriptcentral.com/sigmodrecord>

Prior to submission, please read the Editorial Policy on the SIGMOD Record's website:

<https://sigmodrecord.org/sigmod-record-editorial-policy/>

Rada Chirkova

September 2022

Past SIGMOD Record Editors:

Yanlei Diao (2014-2019)	Ioana Manolescu (2009-2013)	Alexandros Labrinidis (2007-2009)
Mario Nascimento (2005-2007)	Ling Liu (2000-2004)	Michael Franklin (1996-2000)
Jennifer Widom (1995-1996)	Arie Segev (1989-1995)	Margaret H. Dunham (1986-1988)
Jon D. Clark (1984-1985)	Thomas J. Cook (1981-1983)	Douglas S. Kerr (1976-1978)
Randall Rustin (1974-1975)	Daniel O'Connell (1971-1973)	Harrison R. Morse (1969)