Guest Editor's Notes

Welcome to the March 2020 issue of the ACM SIGMOD Record!

The new year of 2020 begins with a special issue on the **2019 ACM SIGMOD Research Highlight Award**. This is an award for the database community to showcase a set of research projects that exemplify core database research. In particular, these projects address an important problem, represent a definitive milestone in solving the problem, and have the potential of significant impact. This award also aims to make the selected works widely known in the database community, to our industry partners, and to the broader ACM community.

The award committee and editorial board included Rada Chirkova, Wim Martens, Jun Yang, and Divesh Srivastava. We solicited articles from PODS 2019, SIGMOD 2019, VLDB 2019, ICDE 2019, EDBT 2019, and ICDT 2019, as well as from community nominations. Through a careful review process eight articles were finally selected as 2019 Research Highlights. The authors of each article worked closely with an associate editor to rewrite the article into a compact 8-page format, and improved it to appeal to the broad data management community. In addition, each research highlight is accompanied by a one-page technical perspective written by an expert on the topic presented in the article. The technical perspective provides the reader with an overview of the background, the motivation, and the key innovation of the featured research highlight, as well as its scientific and practical significance.

The 2019 research highlights cover a broad set of topics, including (a) an automatic way of checking for invariant confluence to enable scaling distributed database systems with consistent semantics ("Checking Invariant Confluence, In Whole or In Parts"); (b) a checkpoint and recovery method with the ability to scale throughput linearly on a large multicore server with negligible increase of latency ("Concurrent Prefix Recovery: Performing CPR on a Database"); (c) the computational complexity of regular document spanners, which provide an abstraction for Information Extraction rules ("Constant-Delay Enumeration for Nondeterministic Document Spanners"); (d) a database perspective on the problem of fairness in machine learning ("Database Repair Meets Algorithmic Fairness"); (e) showing the potential of recursive computations on an RDBMS as the backend for large-scale machine learning ("Declarative Recursive Computation on an RDBMS"); (f) a general framework for detecting if all three problems of enumeration, counting and uniform generation are efficiently solvable ("Efficient Logspace Classes for Enumeration, Counting, and Uniform Generation"); (g) the use of core RDMBS techniques to explain the predictions of a deep ML model ("Query Optimization for Faster Deep CNN Explanations"); and (h) an efficient and tamper-proof way to retrieve and use historical data on a blockchain ("Revealing Every Story of Data in Blockchain Systems").

On behalf of the SIGMOD Record Editorial Board, I hope that you enjoy reading the March 2020 issue of the SIGMOD Record!

Divesh Srivastava

March 2020

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/

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)