Editor's Notes

Welcome to the March 2018 issue of the ACM SIGMOD Record!

The new year of 2018 begins with a special issue on the **2017 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 Zack Ives, Jeff Naughton, Wang-Chiew Tan, and Yanlei Diao. We solicited articles from PODS 2017, SIGMOD 2017, VLDB 2017, ICDE 2017, EDBT 2017, and ICDT 2017, as well as from community nominations. Through a careful review process five articles were finally selected as 2017 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 our associate editor or an external 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 2017 research highlights cover a broad set of topics, including (a) a new theoretical framework for feature engineering for programming machine-learning solutions over a database ("A Relational Framework for Classifier Engineering"); (b) a parallel graph processing system that employs a simple, intuitive programming model and a principled approach based on fixpoint computation which enables database-style optimization ("From Think Parallel to Think Sequential"); (c) a scalable linear algebra system built on top of a parallel relational database system ("Scalable Linear Algebra on a Relational Database System"); (d) an entity matching system that overcomes limitations of existing solutions by considering the requirements for building an end-to-end system ("Magellan: Toward Building Entity Matching Management Systems"); (e) a new approach to helping the user understand answers of natural language queries, i.e., giving an explanation of how and why each answer exists ("Natural Language Explanations for Query Results").

On behalf of the SIGMOD Record Editorial Board, I hope that you enjoy reading the March 2018 issue of the SIGMOD Record!

Your submissions to the SIGMOD Record are welcome via the submission site:
http://sigmod.hosting.acm.org/record

Prior to submission, please read the Editorial Policy on the website of the SIGMOD Record:
http://sigmod.org/sigmodrecord/authors/

Yanlei Diao

March 2018

Past SIGMOD Record Editors:

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