

## Editor's Notes

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

This issue starts with two articles in the Database Principles column. The first article, by Barceló, Kostylev, Monet, Pérez, Reutter, and Silva, surveys recent results concerning architectures of graph neural networks (GNNs) in terms of their ability to classify nodes over graphs. GNNs have recently been proven to be very efficient in many applications, but their theoretical properties are not yet well understood. The work described in the article contributes to better understanding of GNNs, in particular of their power to express node classifiers in graphs. The formal results outlined in the article connect the expressive power of GNNs to unary logical formulas, thus bridging the gap between structure-aware machine-learning architectures, on the one hand, and classic database-query formalisms, on the other hand. The authors also report on experimental corroborations of their results, with the code available online, and discuss open problems and future work in the area.

The second article in the Database Principles column, by Schwentick, Vortmeier, and Zeume, focuses on the problem of dynamic query maintenance, that is of whether query answers can be maintained in response to changes in the database data, by using first-order update rules and potentially auxiliary data. The authors study the problem from the perspective of dynamic complexity theory, and present results centering on the reachability (transitive-closure) query in graphs. The exposition starts from the simplest case and then progresses in a clear sequence of steps each building on previous steps. The article outlines useful techniques, as well as positive and impossibility results, and also discusses implications for regular-path and other types of queries. The authors propose guidelines for determining whether a given query can be dynamically maintained using first-order update rules. The article also presents open problems and provides a discussion of related and further work.

The Vision column presents an article by Amer-Yahia and colleagues on ways to make AI machines work in Future of Work (FoW) scenarios. AI systems are increasingly used for the benefit of humans, and this article focuses specifically on using AI systems to enable human work in both physical and virtual workplaces. Bringing humans to the frontier of FoW would contribute to increasing their trust in AI systems. In the process, human perception could shift to using such systems as a source of self-improvement and better work performance, thus positively shaping national and societal outputs. To make this happen, FoW platforms need to be redesigned, and human workers should be encouraged to take on more supervisory roles, which would allow them to provide corrective feedback to AI systems. The article outlines intellectual challenges that need to be addressed to achieve this vision, including the imperative to capture human capabilities, as well as declarative specification of job-related and workforce-related requirements. The authors also map the intellectual challenges to data-management areas, and use this perspective to review related work.

The Surveys column features an article by Jandre, Diirr, and Braganholo that studies the types of provenance solutions that are available in software tools designed to enable collaborative in-silico research. Collaboration is essential in science, and the emergence of accessible computers and computer networks over the past decades has allowed long-distance collaboration. In fact, it has also increased the number of scientific experiments conducted in silico. The various data and metadata that are collected about objects and activities encountered during in-silico experiments logically belong in provenance databases. The article formulates two main provenance-related challenges in collaborative in-silico experiments and the two corresponding research questions, and provides a taxonomy for, as well as an extended comparison of, state-of-the-art approaches and

provenance-aware models that are available for conducting such experiments. The authors provide literature-based answers to the two research questions, and discuss further challenges and opportunities based on the gaps identified in the survey. The findings presented in the article generate insights that may be useful for researchers interested in the area.

The Distinguished Profiles column features Susan Davidson, professor at the University of Pennsylvania. Sue is an ACM Fellow, a Corresponding Fellow of the Royal Society of Edinburgh, and the recipient of the 2017 IEEE Technical Committee on Data Engineering Impact Award. Her Ph.D. is from Princeton University. In this interview, Sue talks about her research on data provenance and its connections to the problems of data citation and of fake news. She shares how she got interested in bioinformatics and computing, and outlines ideas for doing research that helps domain science or industry. Sue also provides insights on a range of other topics, including what CRA accomplishments she is most proud of, as well as strategies for engaging more women in computer science, including promotion of undergraduate research and sense of community. She talks about balancing work and family, discusses what she would do if she magically had extra time, and gives advice for fledgling and mid-career database researchers.

This issue features a report by Kondylakis, Stefanidis, Rao, and Parry on the outcomes of the Second International Workshop on Semantic Web Meets Health Data Management (SWH 2019). The workshop took place in Auckland, New Zealand in conjunction with the 18<sup>th</sup> International Semantic Web Conference (ISWC 2019). The SWH workshop aimed to bring together an interdisciplinary audience, to discuss challenges in healthcare data management and to propose novel and practical solutions for next-generation data-driven healthcare systems. The article summarizes the outcomes of the workshop and outlines key observations and emerging research directions.

The issue closes with a SIGMOD Executive Committee statement on racism, and a second-round call for SIGMOD 2021 research papers.

On behalf of the SIGMOD Record Editorial board, I hope that you enjoy reading the June 2020 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

June 2020

#### 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)