

Distributed and Parallel Databases

~Special Issue Call for Papers~

“In-Database Analytics”

GUEST EDITORS:

Dan Olteanu (dan.olteanu@cs.ox.ac.uk), University of Oxford

Florin Rusu (frusu@ucmerced.edu), University of California Merced

Distributed and Parallel Databases invites papers for a special issue on In-Database Analytics.

Recent years have witnessed a sustained effort towards big data analytics, such as regression, classification, and recommendation, on massive datasets of varied types and formats. Although most of the methods to perform these tasks have been studied before in machine learning and data mining, scalability and generality have risen as primary challenges not addressed before. As a result, many libraries, frameworks, and platforms have been recently developed to provide scalable support for distributed and parallel statistical analytics. While mostly scalable, several of the newly proposed approaches fall short on generality. In-database solutions seek to reuse and extend the query language and processing infrastructure of a database system with generic constructs and operators required to perform these considerably more complicated tasks. They are generic, but not necessarily scalable.

This special journal issue focuses on conceptual, systems, and algorithmic aspects of large-scale data analytics, both inside and outside the database. Since a large body of work already addresses several facets of this problem, we are seeking contributions that push the state-of-the-art towards novel data processing paradigms for large-scale analytics, as well as those that provide a deep understanding and assessment of the status quo. The main goal of the special issue is to provide a snapshot of the state of the art and a useful point of reference for research in large-scale database analytics.

This special issue seeks articles describing significant research contributions in the domain of large-scale database analytics. Areas of interest include, but are not limited to, the following:

- Data processing architectures for large-scale data analytics
- Approaches for analytics model management and feature engineering
- Data models and query languages for improving or enabling large-scale analytics, including semi- and unstructured models, support for uncertain data, and linear algebra operations
- Methodologies for benchmarking, performance evaluation, and testing data analytics systems
- Exploitation of modern hardware to support analytics
- Processing and optimization techniques focused on data analytics
- Real-world analytics applications, including scientific processing, security, and web analytics

IMPORTANT DATES:

- **Submission deadline: February 1, 2017** (articles will be handled as soon as they are received)
- **First author notification: May 1, 2017**
- **Revised articles due: July 1, 2017**
- **Final author notification: August 15, 2017**

PAPER SUBMISSION:

- Authors are encouraged to submit high-quality, original work that has neither appeared in, nor is under consideration by, other journals.
- All papers will be reviewed following standard reviewing procedures for the Journal.
- Papers must be prepared in accordance with the Journal guidelines: www.springer.com/10619
- Submit manuscripts to: <http://DAPD.edmgr.com>. Choose “In-Database Analytics” as the article type.

Distributed and Parallel Databases www.Springer.com/10619

Amit P. Sheth, Wright State University, USA, Editor-in-Chief

Divyakant Agrawal, University of California at Santa Barbara, USA, Editor-in-Chief

Published by Springer.

