

Editorial

This special issue of the fourteenth volume of the Journal of Information and Data Management (JIDM) comprises extended versions of the three best short papers published in the proceedings of the Brazilian Symposium on Databases (SBBDB) in 2022. SBBDB is the main Brazilian conference and the largest in Latin America on data science and big data. The event covers issues in data management, database and information systems research, data science, big data, and other related topics, considering data as a technological cornerstone of emerging applications. Short papers present initial results or ongoing work in these topics of interest. This issue consists of significantly extended versions of three works selected as the best short papers at the conference.

This special issue starts with the article entitled “*Set Similarity Joins on Heterogeneous Clusters*”, authored by Larissa Silva and Leonardo Ribeiro. Their article presents an approach to evaluating set similarity joins on a heterogeneous cluster of compute nodes equipped with CPU and GPU and a cost model to distribute the workload between these processors, integrating two algorithms, one distributed and the other parallel, in a coprocessing fashion. Experimental results show that the proposal is efficient, scalable, and outperforms previous work.

The second article is “*Mining Temporal Rules from Heterogeneous Multivariate Time Series*”, by Eliane Karasawa and Elaine Sousa. The paper presents TRUMiner (Temporal RULES Miner), an algorithm to mine temporal rules from multivariate time series considering pairs of variables. It provides extended multivariate temporal rules that point to the occurrence of the mined patterns in the original time series, can be used with different discretization methods, and is robust to missing data and heterogeneous time series.

The third article, “*ORBITER: a Lightweight Framework for Automatic Deployment of Big Data Applications on Serverless Architectures*”, by authors João Loureiro, José Maria Monteiro, Marcos Bedo, and Daniel de Oliveira, introduces a lightweight framework for deploying big data applications in a serverless architecture. ORBITER follows the Function-as-a-Service (FaaS) model, relying on open-source tools. The article also shows the results of deploying two practical applications in the Microsoft Azure cloud using the approach, which presented acceptable overhead.

We want to thank everyone who contributed to this Special Edition, particularly reviewers for their insightful comments and authors for their contributions and hard work in preparing their manuscripts.

We hope you enjoy reading this JIDM Special Issue.

Daniel S. Kaster
Special Issue Editor

Daniel de Oliveira
Editor-in-Chief

Ticiano Linhares
Associate Editor