

Table of Contents

Invited Paper

1. Finding Prominent Clusters in VAT Family of Algorithms.....1
Kartik Vishal Deshpande, Osmar R Zaiane, Dheeraj Kumar

Responsible AI & Learning

2. BALM: Bias-Aware Generation for Large Language Models.....16
Dorleon Ginel, Shirin Shujaa
3. FAIRTOPIA: A Multi-Agent Guardianship Framework for Disrupting Unfair AI Pipelines 31
Athena Vakali, Ilias Dimitriadis, Sofia Vei
4. ACAT: A Collaborative Platform for Efficient Aspect-Based Sentiment Dataset Annotation 38
Ana-Maria Luisa Mocanu, Ciprian-Octavian Truicã, Elena-Simona Apostol
5. Calibration improves detection of mislabeled examples44
Ilies Chibane, Thomas George, Pierre Nodet, Vincent Lemaire
6. Q-Matrix Learning Based on Graph Attention Neural Network for Knowledge Tracing 61
Wei Su, Lei Liu, Chuan Cai, Luna Zhang, Yongna Yuan, Shenglin Xu
7. Q-dEclat: A Vertical Quantitative Frequent Pattern Mining Algorithm69
Carson Leung, Adam G.M. Pazdor

Data Lakes & Lakehouses

8. Hybrid Table Annotation in Data Lakes.....75
Nassima Kaid, Zoubida Kedad, Stephane Lopes
9. Context-aware Data Quality Management in Data Lakes90
José de León, Flavia Serra, Adriana Marotta
10. Discovering Relationships in Data Lakes Using Large Language Models: An Industrial Case 105
Ahlame Diouan, Eric Ferey, Sabine Loudcher, Jerome Darmont
11. LLM-Assisted Metadata Query Answering on Data Warehouses120
Matteo Francia, Enrico Gallinucci, Matteo Golfarelli, Manuele Pasini, Stefano Rizzi
12. SmartIngest: A Unified Framework for Adaptive Data Ingestion and Management in Lakehouse Architectures.....135
Ahmed A. Harby, Farhana Zulkernine

Time Series & Machine Learning

13. MoE-TFT: Mixture-of-Experts Enhanced Temporal Fusion Transformer for Time Series Forecasting 151
Saifullah Burero, Anton Dignoes, Michele Linardi, Johann Gamper
14. CODA: A Coordinate-Driven Autoencoder for Robust Multivariate Time Series Anomaly Detection 157
Pierre Lotte, Andre Peninou, Olivier Teste
15. Topologically-Constrained Any-Variate Time-Series Foundation Models for Twinning of Continuous Industrial Processes172
Michael Mayr, Georgios C. Chasparis

16. ConvTransAKD: Multi-Level Wasserstein Alignment for Efficient Knowledge Distillation in Sensor-Based Activity Recognition..... 180
Thi Hong Vuong, Tung Doan, Atsuhiko Takasu
17. ANNOYing HDBSCAN for clustering 186
Ben M. Clark, Ginelle K. Elias, Carson K. Leung, Adam G.M. Pazdor, Thamira M. Randeniya, Xavier J. Schneider
18. Evaluating Time-Dependent Risk Trajectories through Class–Cluster Alignment 201
Sergej Korlakov, Boris Thome, Stefan Conrad

Data Integration & Quality

19. KAN-ER: Kolmogorov-Arnold Networks-based Entity Resolution in Data Lakes 209
Lamisse Bouabdelli, Fatma Abdelhedi, Slimane Hammoudi, Allel Hadjali
20. Leveraging Tattoo Description Classification in Historical Record Linkage ... 224
Eric Sartor, Jeremy Foxcroft, Luiza Antonie
21. ASOS-CRM: Automated Semantic Scoping for CIDOC CRM Population..... 230
Ali Hariri, Stéphane Jean, Mickaël Baron
22. Relational Database Data Lineage Ontology..... 238
Jakub Dutkiewicz, Pawel Misiorek, Robert Wrembel
23. Data Profiling for Change Rules 253
Nishtha Sharma, Fei Chiang
24. Causally-constrained Power Indices: Predicate Attribution for Aggregate Query Explanations 268
Pouya Khani, Yuchen Zhang, Ira Assent

Database & Graph Systems

25. Database Tuning via Distributional Reinforcement Learning for Hybrid Transactional/Analytical Processing Workload 283
Md Rakibul Hasan, Xiaoying Wu, Dimitri Theodoratos
26. How Good are Learned and Hybrid Join Order Strategies from Energy Perspective? 298
Khadija Bennamane, Cheikh Salmi, Ladjel Bellatreche
27. GPU-Accelerated Dependency Graph Construction and Conflict Analysis for Preventing Read-Only Anomalies 313
Reo Chiyomaru, Takamitsu Shioi, Qiong Chang, Jun Miyazaki
28. Evaluating Sustainability in Graph Intelligence 328
Pierre-Paul Cavallera, Landy Andriamampianina, Moncef Garouani, Franck Ravat, Jiefu Song, Nathalie Vallès-Parlangeau
29. A Verifiable Query Processing for Distributed Property Graphs 343
Tomoka Tegawa, Toshiyuki Amagasa

Author Index 349