

1. Record Nr.	UNISA996485665503316
Titolo	Advances in databases and information systems : 26th European conference, ADBIS 2022, Turin, Italy, September 5-8, 2022, proceedings / / Silvia Chiusano, Tania Cerquitelli, Robert Wrembel (editors)
Pubbl/distr/stampa	Cham, Switzerland : , : Springer, , [2022] ©2022
ISBN	3-031-15740-0
Descrizione fisica	1 online resource (419 pages)
Collana	Lecture notes in computer science ; ; Volume 13389
Disciplina	005.74
Soggetti	Database management Information technology
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di bibliografia	Includes bibliographical references and index.
Nota di contenuto	Intro -- Preface -- Organization -- Abstracts of the Keynote Talks -- Toward AI-Powered Data-Driven Education -- Insider Stories: Analyzing Stress, Depression, and Staff Welfare at Major US Companies from Online Reviews -- Tensor Query Processing: Neural Network to speed up Databases and Classical ML! -- Understanding and Rewiring Cities and Societies: A Computational Social Science Perspective -- Contents -- Keynote Talk and Tutorials -- Understanding and Rewiring Cities -- 1 Introduction -- 2 Neighborhoods' Characteristics and Urban Vitality -- 3 How Safety Perception Influences Vitality -- 4 The Interplay of Neighborhoods' Socio-Economic Conditions, Urban Environment, People Behaviors, and Crime -- 5 The Impact of COVID-19 Pandemic on Human Behavior -- 6 Looking Ahead -- References -- AI Approaches in Processing and Using Data in Personalized Medicine -- 1 Introduction -- 2 Different Sources of Patients' Medical Big Data - Collection and Processing -- 3 Emergent Artificial Intelligence Approaches for Supporting Quality Medical Decisions -- 4 Medical Decision Support Systems -- 4.1 Smart Ambient Intelligent Living Environments -- 4.2 Intelligent System for Supporting Cancer Patients -- 5 Conclusion -- References -- Explainable, Interpretable, Trustworthy, Responsible,

Ethical, Fair, Verifiable AI... What's Next? -- 1 Introduction -- 2 Explainability -- 3 Interpretability -- 4 The Problem of Bias -- 5 Fairness -- 6 Verification -- 7 Accountability -- 8 Conclusions -- References -- OLAP and NoSQL: Happily Ever After -- 1 Introduction and Motivation -- 2 Schema-on-Read Approaches -- 2.1 Graph OLAP -- 2.2 Approximate OLAP -- 3 Schema-on-Write Approaches -- 3.1 Mono-Model Approaches -- 3.2 Multi-model Approaches -- 4 Conclusion -- References -- What's New in Temporal Databases? -- 1 Introduction -- 2 Query Processing of Primitive Operators. 2.1 Temporal Selection -- 2.2 Temporal Joins -- 3 New Directions in Models and Semantics -- 4 Systems -- 5 Conclusion and Future Directions -- References -- Graph Processing -- An Algebra for Path Manipulation in Graph Databases -- 1 Introduction -- 2 Related Work -- 3 Data Model -- 4 Path Algebra -- 4.1 Operations over Paths -- 4.2 Operations over Sets of Paths -- 4.3 Examples of Queries -- 5 Properties -- 6 Use Cases -- 6.1 Transportation -- 6.2 Proteins -- 7 Discussion -- 7.1 Labels and Repeated Results -- 7.2 Notions of Compatibility -- 7.3 Paths Queries Across Multiple Graphs -- References -- Road Network Graph Representation for Traffic Analysis and Routing -- 1 Introduction -- 2 Related Work -- 3 Road Network Graph Modelling -- 3.1 Primal Graph -- 3.2 Dual Graph -- 4 Generation of Traffic Data -- 5 Traffic Data Integration -- 6 Routing -- 6.1 Managing Road Closures -- 7 Road Network Analysis -- 8 Conclusion -- References -- Parallel Discovery of Top-k Weighted Motifs in Large Graphs -- 1 Introduction -- 2 Related Research -- 3 Background -- 4 The Proposed Approach -- 4.1 Correctness -- 4.2 Extension to 4-Cliques -- 5 Performance Evaluation -- 6 Conclusions -- References -- A Data Quality Framework for Graph-Based Virtual Data Integration Systems -- 1 Introduction -- 2 Related Work -- 3 Preliminaries -- 3.1 Running Example -- 3.2 Formal Background -- 4 Managing Data Quality in Virtual Data Integration -- 4.1 DC Generation and Graph-Based Representation -- 4.2 Global DC Management -- 5 Global Query Rewriting with Global DCs -- 5.1 Query with DCs -- 6 Validation -- 7 Conclusions and Future Work -- References -- Time Series and Data Streams -- Generating Comparative Explanations of Financial Time Series -- 1 Introduction -- 2 Literature Review -- 3 Data Overview -- 4 Financial Data Summarizer -- 5 Experiments -- 5.1 Intrinsic Evaluation. 5.2 Extrinsic Summary Validation -- 6 Conclusions and Future Works -- References -- Summarizing Edge-Device Data via Core Items -- 1 Introduction -- 2 Related Work -- 3 Preliminaries -- 3.1 Apache Storm -- 3.2 Stream Summarization -- 4 Algorithms for Core Item Sets Computation -- 4.1 SoftSieving -- 5 Experimental Evaluation -- 6 Conclusion -- References -- Parallel Techniques for Variable Size Segmentation of Time Series Datasets -- 1 Introduction -- 2 Problem Definition and Background -- 2.1 SAX Representation -- 2.2 Similarity Queries -- 2.3 Problem Statement -- 3 Adaptive SAX Based on the Representation's Sum of Squared Errors (ASAX_SSE) -- 3.1 Sum of Squared Errors (SSE) -- 3.2 SSE of PAA Representation Considering One Segment (LSSE) -- 3.3 SSE of PAA Representation Considering All Segments (GSSE) -- 3.4 Variable-Size Segmentation Based on SSE Measurement -- 3.5 Lower Bounding of the Similarity Measure -- 4 Parallel Versions of ASAX_SSE -- 4.1 Parallelization on Data -- 4.2 Parallelization on Segments -- 5 Experiments -- 5.1 Setup -- 5.2 Precision of k-Nearest Neighbor Search -- 5.3 Scalability -- 6 Related Work -- 7 Conclusion -- References -- On Line Analytical Processing -- ORTree: Tuning Diversified Similarity Queries by Means of Data Partitioning -- 1 Introduction -- 2 Background -- 2.1 Range-Tree --

2.2 MAM - Omni-Technique -- 2.3 The Diversity Problem -- 2.4 Diversity Algorithms -- 2.5 Candidate Selection -- 3 Methodology -- 4 Experiments -- 4.1 Index Creation Time -- 4.2 Quality Experiments -- 4.3 Number of Elements Retrieved -- 4.4 Query Time Evaluation -- 5 Conclusions and Future Work -- References -- A Knowledge-Based Approach to Support Analytic Query Answering in Semantic Data Lakes*-12pt -- 1 Introduction -- 2 Case Study: Azure COVID-19 Data Lake -- 3 Semantic Data Lake: Data Model -- 3.1 Metadata Layer -- 3.2 Knowledge Layer.

4 Integration and Mapping Discovery -- 5 Query Answering -- 6 Evaluation -- 7 Conclusion -- References -- Insight-Based Vocalization of OLAP Sessions -- 1 Introduction -- 2 Related Work -- 3 Overview -- 4 Formal Background -- 5 The Vocalization Process -- 5.1 Insight Generation -- 5.2 Insight Selection -- 5.3 Vocalization -- 6 A Closer Look at the VOOL Modules -- 7 Evaluation and Conclusion -- References -- Advanced Querying -- Querying Temporal Anomalies in Healthcare Information Systems and Beyond -- 1 Introduction -- 2 Related Work -- 3 Temporal Anomalies -- 3.1 Preliminaries -- 3.2 Definition of Temporal Anomalies -- 3.3 Anomaly Labelling and Retrieval -- 4 SQL Implementations -- 4.1 Unfold/Fold -- 4.2 Unfold/Fold Join Filtered -- 4.3 Window Function -- 5 Experimental Evaluation -- 5.1 Setup and Dataset -- 5.2 Results -- 6 Conclusion and Future Work -- References -- Soft Spatial Querying on JSON Data Sets -- 1 Introduction -- 2 Background -- 2.1 Brief Introduction to Fuzzy Sets -- 2.2 Related Work -- 2.3 The J-CO Framework -- 3 Case Study and J-CO-QL+ Script -- 3.1 Case Study -- 3.2 Definition of Fuzzy Concepts -- 3.3 Processing a GeoJSON Document -- 3.4 Soft Spatial Querying -- 4 Conclusions -- References -- Maximum Range-Sum for Dynamically Occurring Objects with Decaying Weights -- 1 Introduction -- 2 Preliminaries -- 3 Approximate Solution to DDW-MaxRS -- 3.1 Properties of the Approximate Solution -- 3.2 Memory-Efficient Approximate Solution -- 4 Generalization -- 5 Experimental Study -- 6 Related Work -- 7 Conclusions and Future Work -- References -- Performance -- Storage Management with Multi-Version Partitioned B-Trees -- 1 Introduction -- 2 Architecture of Multi-Version Partitioned B-Trees -- 3 Cached Partition: Stop Re-writing Valid Data -- 4 Garbage Collection and Space Reclamation -- 5 Experimental Evaluation -- 6 Conclusion.

References -- Generalization Aware Compression of Molecular Trajectories -- 1 Introduction -- 2 Background -- 3 Methodology -- 4 Experiments -- 5 Conclusions and Future Work -- References -- Analysing Workload Trends for Boosting Triple Stores Performance -- 1 Introduction -- 2 Background and Related Work -- 2.1 RDF Graph and Queries Processing -- 2.2 Indexing and Cache -- 2.3 Replication -- 2.4 Workload Adaption -- 2.5 Heat Query -- 3 Heat Item Gets Colder -- 3.1 Smoothing Methods -- 3.2 Exponential Smoothing -- 3.3 Holt-Winters' Additive Method -- 4 Experimental Evaluation -- 4.1 Exponential Smoothing -- 4.2 Hot Data Parts -- 4.3 Seasonality Factors and Capacity -- 5 Conclusion and Future Work -- References -- Machine Learning -- Comparison of Models Built Using AutoML and Data Fusion -- 1 Introduction -- 2 Related Work -- 3 Problem Statement -- 4 Proposed Approach -- 4.1 Datasets Description -- 4.2 Data Fusion Process -- 4.3 Experimental Design -- 4.4 Performance Evaluation -- 5 Conclusion and Future Work -- References -- Dimensional Data KNN-Based Imputation -- 1 Introduction -- 2 Related Work -- 3 DW Dimension -- 4 Distance Between Dimension Instances -- 4.1 Attribute Distance -- 4.2 Hierarchy Level Instance Distance -- 4.3 Hierarchy Instance Distance -- 4.4 Dimension Instance

Weight -- 5 H-OLAPKNN Imputation -- 5.1 H-OLAPKNN Overview --
5.2 Imputation for Parameters by OLAPKNN -- 6 Experimental
Assessment -- 6.1 Technical Environment and Datasets -- 6.2
Experimental Methodology -- 6.3 Results and Analysis -- 7 Conclusion
and Future Work -- References -- Feature Ranking from Random Forest
Through Complex Network's Centrality Measures*2mm -- 1
Introduction -- 2 Related Work -- 2.1 Feature Selection and Ranking --
2.2 Random Forests -- 2.3 Complex Networks -- 3 From Trees in a
Random Forest to a Complex Network -- 4 Experimental Setup.
5 Results and Discussion.
