

1. Record Nr.	UNISA996466176303316
Titolo	Advances in Spatial and Temporal Databases [[electronic resource]] : 15th International Symposium, SSTD 2017, Arlington, VA, USA, August 21 – 23, 2017, Proceedings // edited by Michael Gertz, Matthias Renz, Xiaofang Zhou, Erik Hoel, Wei-Shinn Ku, Agnes Voisard, Chengyang Zhang, Haiquan Chen, Liang Tang, Yan Huang, Chang-Tien Lu, Siva Ravada
Pubbl/distr/stampa	Cham : , : Springer International Publishing : , : Imprint : Springer, , 2017
ISBN	3-319-64367-3
Edizione	[1st ed. 2017.]
Descrizione fisica	1 online resource (XIV, 454 p. 206 illus.)
Collana	Information Systems and Applications, incl. Internet/Web, and HCI ; ; 10411
Disciplina	005.74
Soggetti	Database management Artificial intelligence Computer science—Mathematics Data mining Database Management Artificial Intelligence Discrete Mathematics in Computer Science Data Mining and Knowledge Discovery
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 -- Contents -- Routing and Trajectories -- Multi-user Itinerary Planning for Optimal Group Preference -- 1 Introduction -- 2 Problem Definition and Preliminaries -- 3 Proposed Solutions -- 3.1 Meeting Graph and Node Profit -- 3.2 Greedy Itinerary Construction -- 3.3 Optimal Itinerary Construction -- 3.4 Acceleration via Graph Compression -- 4 Related Work -- 5 Experiments -- 5.1 Experiment Design -- 5.2 Experimental Results -- 6 Discussions and Conclusion -- References -- Hybrid Best-First Greedy Search for Orienteering with Category Constraints -- 1 Introduction -- 2 Related Work -- 3 Problem Formalization -- 4 Best-

First Search Strategy -- 4.1 Potential Score -- 4.2 Our Algorithm -- 4.3 Further Optimizations -- 5 Approximation Algorithms -- 5.1 Bounding the Score -- 5.2 Bounding the Run Time -- 6 Properties and Bounds -- 6.1 Correctness of Pruning -- 6.2 Lower Bounding the Score -- 6.3 Upper Bounding the Run Time -- 7 Experimental Evaluation -- 7.1 Data Sets -- 7.2 Effects of Parameters -- 7.3 Comparison with Competitors -- 8 Conclusion and Future Work -- References -- On Privacy in Spatio-Temporal Data: User Identification Using Microblog Data -- 1 Introduction -- 2 Related Work -- 2.1 User Identification -- 2.2 User Linkage -- 2.3 Spatial Privacy -- 3 Problem Definition -- 4 Trajectory Based User Identification -- 4.1 Trace Profile Modeling -- 4.2 Set Descriptors -- 4.3 Transition Descriptors -- 4.4 Classification -- 4.5 User Linkage -- 5 Experimental Evaluation -- 5.1 Accuracy Using Set Descriptors -- 5.2 Accuracy Using Frequent Transitions -- 5.3 Accuracy for Different Observation Counts -- 5.4 User Linkage Between Different Social Networks -- 5.5 Scalability -- 6 Conclusions -- References -- Big Spatial Data -- Sphinx: Empowering Impala for Efficient Execution of SQL Queries on Big Spatial Data. 1 Introduction -- 2 Background on Impala -- 3 Architecture -- 4 Query Parser -- 4.1 Geometry Data Type -- 4.2 Spatial Functions -- 4.3 Spatial Operations -- 4.4 Spatial Indexing -- 5 Spatial Indexing -- 5.1 Index Construction in Sphinx -- 5.2 Importing SpatialHadoop Indexes -- 6 Query Planner -- 6.1 Range Query Plans -- 6.2 Spatial Join Plans -- 7 Query Executor -- 7.1 R-tree Scanner -- 7.2 Spatial Join Operator -- 8 Experiments -- 8.1 Index Construction -- 8.2 Range Query -- 8.3 Spatial Join -- 9 Related Work -- 10 Conclusion -- References -- ST-Hadoop: A MapReduce Framework for Spatio-Temporal Data -- 1 Introduction -- 2 Related Work -- 3 ST-Hadoop Architecture -- 4 Language Layer -- 5 Indexing Layer -- 5.1 Concept of Hierarchy -- 5.2 Index Construction -- 5.3 Phase I: Sampling -- 5.4 Phase II: Temporal Slicing -- 5.5 Phase III: Spatial Indexing -- 5.6 Phase IV: Physical Writing -- 6 Operations Layer -- 6.1 Spatio-Temporal Range Query -- 6.2 Spatio-Temporal Join -- 7 Experiments -- 7.1 Spatiotemporal Range Query -- 7.2 Index Construction -- 7.3 Spatiotemporal Join -- 8 Conclusion -- References -- GeoWave: Utilizing Distributed Key-Value Stores for Multidimensional Data -- Abstract -- 1 Introduction -- 2 Related Work -- 3 Contributions -- 3.1 Locality Preservation of Multidimensional Values -- 3.2 Spatial Subsampling for Map Rendering -- 3.3 Managing Data Variety and Complexity -- 3.4 Key-Value Store Parity -- 4 Experimental Evaluation -- 4.1 Locality Preservation Performance -- 4.2 Map Pixel-Based Spatial Subsampling Performance -- 4.3 Differences in Multi-range Scans Among Key-Value Stores -- 5 Conclusions -- References -- Indexing and Aggregation -- Sweeping-Based Temporal Aggregation -- 1 Introduction -- 2 Related Work -- 3 Problem Formalization -- 3.1 Temporal Relations -- 3.2 Temporal Aggregation on Constant Intervals. 3.3 Temporal Aggregation on Fixed Intervals -- 4 Sweeping-Based Temporal Aggregation -- 4.1 Endpoint Index -- 4.2 Temporal Aggregation on Constant Intervals -- 4.3 Temporal Aggregation on Fixed Intervals -- 5 Empirical Evaluation -- 5.1 Environment -- 5.2 Competitors -- 5.3 Test Workloads -- 5.4 Results -- 6 Conclusion -- References -- Indexing the Pickup and Drop-Off Locations of NYC Taxi Trips in PostgreSQL -- Lessons from the Road -- 1 Introduction -- 2 Studied Spatial Database Indexing Schemes -- 2.1 Generalized Search Tree (GiST-Spatial) -- 2.2 Block Range Index (BRIN-Spatial) -- 2.3 Hippo-Spatial -- 3 Experimental Environment -- 4 Studying the Indexing Overhead -- 4.1 Index Size -- 4.2 Index Initialization Time -- 5 Evaluating the Query Response Time -- 5.1 Varying the Spatial Range

Query Selectivity Factor -- 5.2 Varying the Spatial Range Area Size -- 6
Studying the Index Maintenance Overhead -- 6.1 Insertion Time -- 6.2
Deletion Time -- 6.3 Hybrid Workload Performance -- 7 Summary of
Results -- 8 Key Insights and Learned Lessons -- References --
Towards Spatially- and Category-Wise k-Diverse Nearest Neighbors
Queries -- 1 Introduction -- 2 Related Work -- 3 Problem Definition --
4 Proposed Solutions -- 4.1 Recursive Range Filtering (RRF) -- 4.2 Pair
Graph Solution (PG) -- 4.3 From Conventional Skylines to Linear
Skylines -- 5 Experiments -- 5.1 Spatial Diversity -- 5.2 Categorical
Diversity -- 5.3 Effectiveness and Efficiency -- 6 Conclusion --
References -- Spatio-Temporal Functional Dependencies for Sensor
Data Streams -- 1 Introduction -- 2 Preliminaries -- 3 Database
Modeling for Sensor Data -- 3.1 Granularity Aware Sensor Database --
3.2 Spatio-Temporal Functional Dependency (STFD) -- 3.3 Reasoning
on STFDs -- 3.4 Normalization -- 3.5 Semantic Value Assumption -- 4
Prototype for Sensor Database -- 4.1 Implementation.
4.2 Ongoing Experiments -- 5 Related Work -- 6 Conclusion --
References -- Recommendation -- Location-Aware Query
Recommendation for Search Engines at Scale -- 1 Introduction -- 2
Preliminaries and Definitions -- 2.1 Query Log -- 2.2 Obtaining
Locations from a Query Log -- 2.3 Query-Flow Graph -- 2.4 Term-
Query-Flow Graph -- 3 Location-Aware Query Recommendation -- 4
Location-Aware PPR -- 4.1 BCA with Online Transition Matrix -- 5
Experimental Evaluation -- 5.1 Dataset -- 5.2 Methodology -- 5.3 User
Study -- 5.4 Effectiveness -- 5.5 Efficiency -- 6 Related Work -- 7
Conclusion -- References -- Top-k Taxi Recommendation in Realtime
Social-Aware Ridesharing Services -- 1 Introduction -- 2 Related Work
-- 2.1 Static Ridesharing -- 2.2 Dynamic Ridesharing -- 2.3 Trust-
Conscious Ridesharing -- 3 Problem Formulation -- 3.1 Definitions --
3.2 Spatial and Social Scores -- 3.3 Solution Overview -- 4 Candidate
Taxis Searching -- 4.1 Edge-Based Candidates Selection -- 4.2 Grid-
Based Candidates Selection -- 5 Taxi Scheduling and Top-k Taxi
Selection -- 5.1 Overall Procedure of Top-k Taxis Selection -- 5.2
Spatial Score Upper Bounds -- 5.3 Time-Dependent Fastest Path
Calculation -- 5.4 Optimal Schedule -- 5.5 Hopping Algorithm -- 6
Experimental Evaluation -- 6.1 Experimental Settings -- 6.2
Experimental Results -- 7 Conclusion -- References -- P-LAG:
Location-Aware Group Recommendation for Passive Users -- 1
Introduction -- 2 Problem Definition -- 3 Vector Extraction -- 3.1
Topic Vector Extraction -- 3.2 Topic Vector Analysis -- 4 Indexing and
Search for P-LAG -- 4.1 Basic R-tree Approach -- 4.2 TAR-tree
Approach: Topic-Aware R-tree -- 4.3 Vector Compression in the TAR-
tree -- 5 Experimental Evaluation -- 5.1 Datasets -- 5.2 Efficiency
Analysis -- 5.3 Effectiveness -- 5.4 Storage Requirements -- 6 Related
Work -- 7 Conclusion -- References -- Data Mining.
Grid-Based Colocation Mining Algorithms on GPU for Big Spatial Event
Data: A Summary of Results -- 1 Introduction -- 2 Problem Statement
-- 2.1 Basic Concepts -- 2.2 Problem Definition -- 3 Proposed
Approach -- 3.1 Algorithm Overview -- 3.2 Cell-Aggregate-Based
Upper Bound Filter -- 3.3 Refinement Algorithms -- 4 Evaluation --
4.1 Results on Synthetic Data -- 4.2 Results on Real World Dataset -- 5
Discussion -- 6 Conclusion and Future Work -- References --
Detecting Isodistance Hotspots on Spatial Networks: A Summary of
Results -- 1 Introduction -- 2 Problem Statement -- 2.1 Basic
Concepts -- 2.2 Problem Formulation -- 3 BaseNIHD: A Baseline
Algorithm Using Known Algorithmic Refinements -- 4 NPP: An
Algorithm Based on Network Partitioning and Upper-Bound Pruning --
5 Theoretical Analysis -- 6 Case Studies on Real World Crime Data --

6.1 Robberies Occurred in Pinellas County, Florida -- 6.2 Assaults Occurred in Fremont, Washington -- 7 Experimental Evaluation -- 7.1 Experimental Setup -- 7.2 Experimental Results -- 8 Conclusion and Future Work -- References -- Detection and Prediction of Natural Hazards Using Large-Scale Environmental Data -- 1 Introduction -- 1.1 Roadmap -- 2 Framework for Natural Hazards Detection -- 2.1 Environmental Tensor Factorization -- 2.2 Classifying Natural Hazards -- 3 Spatio-Temporal Tensor Sparsification -- 3.1 Algorithmic procedure -- 4 Experimental Evaluation -- 4.1 Global Climate Data -- 4.2 Finding Synthetic Spatio-Temporal Outliers -- 4.3 Finding Natural Hazards on Real Data -- 5 Related Work -- 5.1 Spatio-Temporal Outlier Detection -- 5.2 Classification and Prediction -- 5.3 Tensor Factorization -- 6 Conclusion -- References -- Localization and Spatial Allocation -- FF-SA: Fragmentation-Free Spatial Allocation -- 1 Introduction -- 2 Problem Statement: Fragmentation-Free Spatial Allocation -- 3 Challenges. 4 Related Work and Limitations.

Sommario/riassunto

This book constitutes the refereed proceedings of the 15th International Symposium on Spatial and Temporal Databases, SSTD 2017, held in Arlington, VA, USA, in August 2017. The 19 full papers presented together with 8 demo papers and 5 vision papers were carefully reviewed and selected from 90 submissions. The papers are organized around the current research on concepts, tools, and techniques related to spatial and temporal databases. .
