

2019 IEEE 35th International Conference on Data Engineering (ICDE) ICDE 2019

Table of Contents

Message from the ICDE 2019 Chairs	xxxiv
Organizing Committee	xxxvi
Program Committee	xxxviii

Keynotes

AsterixDB Mid-Flight: A Case Study in Building Systems in Academia	1
<i>Michael J. Carey (University of California, Irvine)</i>	
Data Management at Huawei: Recent Accomplishments and Future Challenges	13
<i>Jianjun Chen (Huawei America Research), Yu Chen (Huawei America Research), Zhibiao Chen (Huawei Technologies Co., Ltd.), Ahmad Ghazal (Huawei America Research), Guoliang Li (Huawei Technologies Co., Ltd.), Sihao Li (Huawei Technologies Co., Ltd.), Weijie Ou (Huawei Technologies Co., Ltd.), Yang Sun (Huawei America Research), Mingyi Zhang (Huawei America Research), and Minqi Zhou (Huawei Technologies Co., Ltd.)</i>	
Building a Broad Knowledge Graph for Products	25
<i>Xin Luna Dong (Amazon Inc)</i>	
Knowledge Graphs and Enterprise AI: The Promise of an Enabling Technology	26
<i>Luigi Bellomarini (Banca d'Italia), Daniele Fakhoury (Universita Tor Vergata), Georg Gottlob (University of Oxford & TU Wien), and Emanuel Sallinger (University of Oxford & TU Wien)</i>	

Panel

Is There a Data Science and Engineering Brain Drain? If So, How Can We Rebalance Them?	38
<i>Jian Pei (JD.com and Simon Fraser University)</i>	

Research Track

Research (1) - Graphs 1: Query Processing 1

Answering Why-Questions for Subgraph Queries in Multi-attributed Graphs	40
<i>Qi Song (Washington State University), Mohammad Hossein Namaki (Washington State University), and Yinghui Wu (Washington State University; Pacific Northwest National Laboratory)</i>	
Enumerating k-Vertex Connected Components in Large Graphs	52
<i>Dong Wen (University of Technology Sydney), Lu Qin (University of Technology Sydney), Ying Zhang (University of Technology Sydney), Lijun Chang (University of Sydney), and Ling Chen (University of Technology Sydney)</i>	
Index-Based Optimal Algorithm for Computing K-Cores in Large Uncertain Graphs	64
<i>Bohua Yang (University of Technology Sydney), Dong Wen (University of Technology Sydney), Lu Qin (University of Technology Sydney), Ying Zhang (University of Technology Sydney), Lijun Chang (The University of Sydney), and Rong-Hua Li (Beijing Institute of Technology)</i>	
Computing a Near-Maximum Independent Set in Dynamic Graphs	76
<i>Weiguo Zheng (Fudan University), Chengzhi Piao (The Chinese University of Hong Kong), Hong Cheng (The Chinese University of Hong Kong), and Jeffrey Xu Yu (The Chinese University of Hong Kong)</i>	
Contextual Community Search Over Large Social Networks	88
<i>Lu Chen (Swinburne University of Technology), Chengfei Liu (Swinburne University of Technology), Kewen Liao (Charles Darwin University), Jianxin Li (Deakin University), and Rui Zhou (Swinburne University of Technology)</i>	

Research (2) - Distributed, Parallel and P2P Data Management

Rima: An RDMA-Accelerated Model-Parallelized Solution to Large-Scale Matrix Factorization	100
<i>Jinkun Geng (Tsinghua University), Dan Li (Tsinghua University), and Shuai Wang (Tsinghua University)</i>	
Accelerating Partial Evaluation in Distributed SPARQL Query Evaluation	112
<i>Peng Peng (Hunan University), Lei Zou (Peking University; Beijing Institute of Big Data Research; National Engineering Laboratory for Big Data Analysis Technology and Application (PKU), China), and Runyu Guan (Hunan University)</i>	
Blockplane: A Global-Scale Byzantizing Middleware	124
<i>Faisal Nawab (University of California) and Mohammad Sadoghi (UC Davis)</i>	
BENU: Distributed Subgraph Enumeration with Backtracking-Based Framework	136
<i>Zhaokang Wang (Nanjing University), Rong Gu (Nanjing University), Weiwei Hu (Nanjing University), Chunfeng Yuan (Nanjing University), and Yihua Huang (Nanjing University)</i>	

Efficient Synchronization of State-Based CRDTs	148
<i>Vitor Enes (HASLab / INESC TEC and Universidade do Minho), Paulo Sérgio Almeida (HASLab / INESC TEC and Universidade do Minho), Carlos Baquero (HASLab / INESC TEC and Universidade do Minho), and João Leitão (NOVA LINCIS, FCT and Universidade NOVA de Lisbon)</i>	

Research (3) - Data Cleaning

Learning Individual Models for Imputation	160
<i>Aoqian Zhang (Tsinghua University), Shaoxu Song (Tsinghua University), Yu Sun (Tsinghua University), and Jianmin Wang (Tsinghua University)</i>	
CurrentClean: Spatio-Temporal Cleaning of Stale Data	172
<i>Zheng Zheng (McMaster University), Mostafa Milani (McMaster University), and Fei Chiang (McMaster University)</i>	
Fine-Grained Provenance for Matching & ETL	184
<i>Nan Zheng (University of Pennsylvania), Abdussalam Alawini (University of Illinois, Urbana-Champaign), and Zachary G. Ives (University of Pennsylvania)</i>	
Unsupervised String Transformation Learning for Entity Consolidation	196
<i>Dong Deng (Rutgers University), Wenbo Tao (MIT), Ziawasch Abedjan (TU Berlin), Ahmed Elmagarmid (QCRI HBKU), Ihab F. Ilyas (University of Waterloo), Guoliang Li (Tsinghua University), Samuel Madden (MIT), Mourad Ouzzani (QCRI HBKU), Michael Stonebraker (MIT), and Nan Tang (QCRI HBKU)</i>	
A Semi-Supervised Framework of Clustering Selection for De-Duplication	208
<i>Shrinu Kushagra (University of Waterloo), Hemant Saxena (University of Waterloo), Ihab F. Ilyas (University of Waterloo), and Shai Ben-David (University of Waterloo)</i>	

Research (4) - Graphs 2: Query Processing and Indexing

Scaling Up Subgraph Query Processing with Efficient Subgraph Matching	220
<i>Shixuan Sun (Hong Kong University of Science and Technology) and Qiong Luo (Hong Kong University of Science and Technology)</i>	
Efficient Parallel Subgraph Enumeration on a Single Machine	232
<i>Shixuan Sun (Hong Kong University of Science and Technology), Yulin Che (Hong Kong University of Science and Technology), Lipeng Wang (Hong Kong University of Science and Technology), and Qiong Luo (Hong Kong University of Science and Technology)</i>	
Fast Dual Simulation Processing of Graph Database Queries	244
<i>Stephan Mennicke (Institute for Information Systems, TU Braunschweig), Jan-Christoph Kalo (Institute for Information Systems, TU Braunschweig), Denis Nagel (Institute for Information Systems, TU Braunschweig), Hermann Kroll (Institute for Information Systems, TU Braunschweig), and Wolf-Tilo Balke (Institute for Information Systems, TU Braunschweig)</i>	

Efficient and Incremental Clustering Algorithms on Star-Schema Heterogeneous Graphs	256
<i>Lu Chen (Zhejiang University), Yunjun Gao (Zhejiang University), Yuanliang Zhang (Zhejiang University), Christian S. Jensen (Aalborg University), and Bolong Zheng (Aalborg University)</i>	
G*-Tree: An Efficient Spatial Index on Road Networks	268
<i>Zijian Li (Hong Kong University of Science and Technology), Lei Chen (Hong Kong University of Science and Technology), and Yue Wang (Hong Kong University of Science and Technology)</i>	

Research (5) - Recommendation and Clustering

DBSVEC: Density-Based Clustering Using Support Vector Expansion	280
<i>Zhen Wang (Tsinghua University), Rui Zhang (The University of Melbourne), Jianzhong Qi (The University of Melbourne), and Bo Yuan (Tsinghua University)</i>	
A Joint Context-Aware Embedding for Trip Recommendations	292
<i>Jiayuan He (The University of Melbourne), Jianzhong Qi (The University of Melbourne), and Kotagiri Ramamohanarao (The University of Melbourne)</i>	
AIR: Attentional Intention-Aware Recommender Systems	304
<i>Tong Chen (The University of Queensland), Hongzhi Yin (The University of Queensland), Hongxu Chen (The University of Queensland), Rui Yan (Peking University), Quoc Viet Hung Nguyen (Griffith University), and Xue Li (The University of Queensland and Dalian Neusoft University of Information)</i>	
No, That's Not My Feedback: TV Show Recommendation Using Watchable Interval	316
<i>Kyung-Jae Cho (Hanyang University), Yeon-Chang Lee (Hanyang University), Kyungsik Han (Ajou University), Jaeho Choi (Naver Corporation), and Sang-Wook Kim (Hanyang University)</i>	
Adaptive Wavelet Clustering for Highly Noisy Data	328
<i>Zengjian Chen (Huazhong University of Science and Technology), Jiayi Liu (University of Massachusetts Amherst), Yihe Deng (University of California, Los Angeles), Kun He (Huazhong University of Science and Technology), and John E. Hopcroft (Cornell University)</i>	

Research (6) - Strings, Texts, and Keyword Search

An Efficient Parallel Keyword Search Engine on Knowledge Graphs	338
<i>Yueji Yang (National University of Singapore), Divy Kant Agrawal (University of California at Santa Barbara), H. V. Jagadish (University of Michigan of Ann Arbor), Anthony K. H. Tung (National University of Singapore), and Shuang Wu (National University of Singapore)</i>	
Towards Longitudinal Analytics on Social Media Data	350
<i>Fan Xia (East China Normal University), Bin Yang (Aalborg University), Chengcheng Yu (Shanghai Polytechnic University), Weining Qian (East China Normal University), and Aoying Zhou (East China Normal University)</i>	

LCJoin: Set Containment Join via List Crosscutting	362
<i>Dong Deng (Inception Institute of Artificial Intelligence), Chengcheng Yang (Inception Institute of Artificial Intelligence), Shuo Shang (Inception Institute of Artificial Intelligence), Fan Zhu (Inception Institute of Artificial Intelligence), Li Liu (Inception Institute of Artificial Intelligence), and Ling Shao (Inception Institute of Artificial Intelligence)</i>	
Bridging the Semantic Gap with SQL Query Logs in Natural Language Interfaces to Databases	374
<i>Christopher Baik (University of Michigan), H. V. Jagadish (University of Michigan), and Yunyao Li (IBM Research - Almaden)</i>	
MF-Join: Efficient Fuzzy String Similarity Join with Multi-level Filtering	386
<i>Jin Wang (UCLA), Chunbin Lin (Amazon AWS), and Carlo Zaniolo (UCLA)</i>	

Research (7) - Graphs 3: Social Networks

Finding Temporal Influential Users Over Evolving Social Networks	398
<i>Shixun Huang (RMIT University), Zhifeng Bao (RMIT University), J. Shane Culpepper (RMIT University), and Bang Zhang (Damo Academy Alibaba)</i>	
Seed Selection and Social Coupon Allocation for Redemption Maximization in Online Social Networks	410
<i>Tung-Chun Chang (Research Center for Information Technology Innovation, Academia Sinica, Taiwan), Yishuo Shi (Institute of Information Science, Academia Sinica, Taiwan), De-Nian Yang (Institute of Information Science, Academia Sinica, Taiwan), and Wen-Tsuen Chen (Department of Computer Science, National Tsing Hua University, Taiwan)</i>	
Keyword-Centric Community Search	422
<i>Zhiwei Zhang (Hong Kong Baptist University), Xin Huang (Hong Kong Baptist University), Jianliang Xu (Hong Kong Baptist University), Byron Choi (Hong Kong Baptist University), and Zechao Shang (The University of Chicago)</i>	
Cohesive Group Nearest Neighbor Queries Over Road-Social Networks	434
<i>Fangda Guo (Northeastern University), Ye Yuan (Northeastern University), Guoren Wang (Beijing Institute of Technology), Lei Chen (Hong Kong University of Science and Technology), Xiang Lian (Kent State University), and Zimeng Wang (Northeastern University)</i>	
Maximizing Multifaceted Network Influence	446
<i>Yuchen Li (Singapore Management University), Ju Fan (Renmin University of China), George Ovchinnikov (Skolkovo Institute of Science and Technology), and Panagiotis Karras (Aarhus University)</i>	

Research (8) - Approximation

GB-KMV: An Augmented KMV Sketch for Approximate Containment Similarity Search	458
<i>Yang Yang (University of New South Wales), Ying Zhang (University of Technology Sydney), Wenjie Zhang (University of New South Wales), and Zengfeng Huang (Fudan University)</i>	

ARROW: Approximating Reachability Using Random Walks Over Web-Scale Graphs	470
<i>Neha Sengupta (Indian Institute of Technology Delhi), Amitabha Bagchi (Indian Institute of Technology Delhi), Maya Ramanath (Indian Institute of Technology Delhi), and Srikanta Bedathur (Indian Institute of Technology Delhi)</i>	
Taster: Self-Tuning, Elastic and Online Approximate Query Processing	482
<i>Matthaios Olma (EPFL), Odysseas Papapetrou (TU Eindhoven), Raja Appuswamy (EURECOM), and Anastasia Ailamaki (EPFL)</i>	
An Iterative Scheme for Leverage-Based Approximate Aggregation	494
<i>Shanshan Han (Harbin Institute of Technology), Hongzhi Wang (Harbin Institute of Technology), Jialin Wan (Harbin Institute of Technology), and Jianzhong Li (Harbin Institute of Technology)</i>	
Deletion Propagation for Multiple Key Preserving Conjunctive Queries: Approximations and Complexity	506
<i>Zhipeng Cai (Georgia State University), Dongjing Miao (Harbin Institute of Technology; Georgia State University), and Yingshu Li (Georgia State University)</i>	

Research (9) - Machine Learning and Data Science

Enumerating Minimal Weight Set Covers	518
<i>Zahi Ajami (The Hebrew University of Jerusalem) and Sara Cohen (The Hebrew University of Jerusalem)</i>	
Constraints-Based Explanations of Classifications	530
<i>Daniel Deutch (Tel Aviv University) and Nave Frost (Tel Aviv University)</i>	
KARL: Fast Kernel Aggregation Queries	542
<i>Tsz Nam Chan (Hong Kong Polytechnic University, The University of Hong Kong), Man Lung Yiu (Hong Kong Polytechnic University), and Leong Hou U (University of Macau)</i>	
Assessing and Remediating Coverage for a Given Dataset	554
<i>Abolfazl Asudeh (University of Michigan), Zhongjun Jin (University of Michigan), and H. V. Jagadish (University of Michigan)</i>	

Research (10) - Graphs 4: Mining 1

Social Influence-Based Group Representation Learning for Group Recommendation	566
<i>Hongzhi Yin (The University of Queensland), Qinyong Wang (The University of Queensland), Kai Zheng (University of Electronic Science and Technology of China), Zhixu Li (Soochow University), Jiali Yang (Soochow University), and Xiaofang Zhou (The University of Queensland)</i>	
MIDAS: Finding the Right Web Sources to Fill Knowledge Gaps	578
<i>Xiaolan Wang (University of Massachusetts Amherst), Xin Luna Dong (Amazon Inc.), Yang Li (Google Inc.), and Alexandra Meliou (University of Massachusetts Amherst)</i>	

Exploiting Centrality Information with Graph Convolutions for Network Representation Learning	590
<i>Hongxu Chen (The University of Queensland), Hongzhi Yin (The University of Queensland), Tong Chen (The University of Queensland), Quoc Viet Hung Nguyen (Griffith University), Wen-Chih Peng (National Chiao Tung University), and Xue Li (The University of Queensland; Dalian Neusoft University of Information)</i>	
Route Recommendations on Road Networks for Arbitrary User Preference Functions	602
<i>Pranali Yawalkar (Google) and Sayan Ranu (IIT Delhi)</i>	
NSCaching: Simple and Efficient Negative Sampling for Knowledge Graph Embedding	614
<i>Yongqi Zhang (the Hong Kong University of Science and Technology), Quanming Yao (4Paradigm), Yingxia Shao (Beijing University of Posts and Telecommunications), and Lei Chen (the Hong Kong University of Science and Technology)</i>	

Research (11) - Database Privacy, Security, and Trust

ServeDB: Secure, Verifiable, and Efficient Range Queries on Outsourced Database	626
<i>Songrui Wu (Tsinghua University), Qi Li (Tsinghua University), Guoliang Li (Tsinghua University), Dong Yuan (Tsinghua University), Xingliang Yuan (Monash University), and Cong Wang (City University of Hong Kong)</i>	
Collecting and Analyzing Multidimensional Data with Local Differential Privacy	638
<i>Ning Wang (Ocean University of China), Xiaokui Xiao (National University of Singapore), Yin Yang (Hamad Bin Khalifa University), Jun Zhao (Nanyang Technological University), Siu Cheung Hui (Nanyang Technological University), Hyejin Shin (Samsung Electronics), Junbum Shin (Samsung Electronics), and Ge Yu (Northeastern University)</i>	
Partitioned Data Security on Outsourced Sensitive and Non-Sensitive Data	650
<i>Sharad Mehrotra (University of California, Irvine), Shantanu Sharma (University of California, Irvine), Jeffrey Ullman (Stanford University), and Anurag Mishra (University of California, Irvine)</i>	
SecEQP: A Secure and Efficient Scheme for SkNN Query Problem Over Encrypted Geodata on Cloud	662
<i>Xinyu Lei (Michigan State University), Alex X. Liu (Michigan State University), Rui Li (Dongguan University of Technology), and Guan-Hua Tu (Michigan State University)</i>	
Joins Over Encrypted Data with Fine Granular Security	674
<i>Florian Hahn (SAP SE), Nicolas Loza (SAP SE), and Florian Kerschbaum (University of Waterloo)</i>	

Research (12) - Modern Hardware and In-Memory Database Systems

Column-Oriented Database Acceleration Using FPGAs	686
<i>Satoru Watanabe (Hitachi, Ltd.), Kazuhisa Fujimoto (Hitachi, Ltd.), Yuji Saeki (Hitachi, Ltd.), Yoshifumi Fujikawa (Hitachi, Ltd.), and Hiroshi Yoshino (Hitachi ULSI Systems Co., Ltd.)</i>	

Hardware-Conscious Hash-Joins on GPUs	698
<i>Panagiotis Sioulas (École polytechnique fédérale de Lausanne), Periklis Chrysogelos (École polytechnique fédérale de Lausanne), Manos Karpapothakis (École polytechnique fédérale de Lausanne), Raja Appuswamy (École polytechnique fédérale de Lausanne), and Anastasia Ailamaki (École polytechnique fédérale de Lausanne)</i>	
TuFast: A Lightweight Parallelization Library for Graph Analytics	710
<i>Zechao Shang (University of Chicago), Jeffrey Xu Yu (Chinese University of Hong Kong), and Zhiwei Zhang (Hong Kong Baptist University)</i>	
LDC: A Lower-Level Driven Compaction Method to Optimize SSD-Oriented Key-Value Stores	722
<i>Yunpeng Chai (Renmin University of China), Yanfeng Chai (Renmin University of China), Xin Wang (Tianjin University), Haocheng Wei (Renmin University of China), Ning Bao (Renmin University of China), and Yushi Liang (Renmin University of China)</i>	
No False Negatives: Accepting All Useful Schedules in a Fast Serializable Many-Core System	734
<i>Dominik Durner (Technical University of Munich) and Thomas Neumann (Technical University of Munich)</i>	

Research (13) - Graphs 5: Mining 2

Discovering Maximal Motif Cliques in Large Heterogeneous Information Networks	746
<i>Jiafeng Hu (The University of Hong Kong), Reynold Cheng (The University of Hong Kong), Kevin Chen-Chuan Chang (University of Illinois at Urbana-Champaign), Aravind Sankar (University of Illinois at Urbana-Champaign), Yixiang Fang (The University of Hong Kong), and Brian Y.H. Lam (University of Cambridge)</i>	
REPT: A Streaming Algorithm of Approximating Global and Local Triangle Counts in Parallel	758
<i>Pinghui Wang (Xi'an Jiaotong University), Peng Jia (Xi'an Jiaotong University, China), Yiyang Qi (Xi'an Jiaotong University, China), Yu Sun (Xi'an Jiaotong University, China), Jing Tao (Xi'an Jiaotong University), and Xiaohong Guan (Xi'an Jiaotong University)</i>	
Information Diffusion Prediction via Recurrent Cascades Convolution	770
<i>Xueqin Chen (University of Electronic Science and Technology of China), Fan Zhou (University of Electronic Science and Technology of China), Kunpeng Zhang (University of Maryland, College Park), Goce Trajcevski (Iowa State University), Ting Zhong (University of Electronic Science and Technology of China), and Fengli Zhang (University of Electronic Science and Technology of China)</i>	
Finding Densest Lasting Subgraphs in Dynamic Graphs: A Stochastic Approach	782
<i>Xuanming Liu (University of Massachusetts, Lowell), Tingjian Ge (University of Massachusetts, Lowell), and Yinghui Wu (Washington State University)</i>	
Multicapacity Facility Selection in Networks	794
<i>Alvis Logins (Aarhus University), Panagiotis Karras (Aarhus University), and Christian S. Jensen (Aalborg University)</i>	

Research (14) - Query Processing, Indexing and Optimization

An MBR-Oriented Approach for Efficient Skyline Query Processing	806
<i>Ji Zhang (Auburn University), Wenlu Wang (Auburn University), Xunfei Jiang (Earlham College), Wei-Shinn Ku (Auburn University), and Hua Lu (Aalborg University)</i>	
Dynamic Set kNN Self-Join	818
<i>Daichi Amagata (Osaka University), Takahiro Hara (Osaka University), and Chuan Xiao (Nagoya University)</i>	
Packed Memory Arrays - Rewired	830
<i>Dean De Leo (CWI) and Peter Boncz (CWI)</i>	
GEM ² -Tree: A Gas-Efficient Structure for Authenticated Range Queries in Blockchain	842
<i>Ce Zhang (Hong Kong Baptist University), Cheng Xu (Hong Kong Baptist University), Jianliang Xu (Hong Kong Baptist University), Yuzhe Tang (Syracuse University), and Byron Choi (Hong Kong Baptist University)</i>	
Effective Filters and Linear Time Verification for Tree Similarity Joins	854
<i>Thomas Hütter (University of Salzburg), Mateusz Pawlik (University of Salzburg), Robert Löschinger (University of Salzburg), and Nikolaus Augsten (University of Salzburg)</i>	

Research (15) - Temporal and Spatial Data 1

KV-Match: A Subsequence Matching Approach Supporting Normalization and Time Warping	866
<i>Jiaye Wu (Fudan University), Peng Wang (Fudan University), Ningting Pan (Fudan University), Chen Wang (Tsinghua University), Wei Wang (Fudan University), and Jianmin Wang (Tsinghua University)</i>	
Efficient Maximal Spatial Clique Enumeration	878
<i>Chen Zhang (University of New South Wales, Australia), Ying Zhang (University of Technology Sydney, Australia), Wenjie Zhang (University of New South Wales, Australia), Lu Qin (University of Technology Sydney, Australia), and Jianye Yang (Ant Financial Services Group, Alibaba Group, China)</i>	
Cluster-Based Subscription Matching for Geo-Textual Data Streams	890
<i>Lisi Chen (Inception Institute of Artificial Intelligence (IIAI), UAE), Shuo Shang (UESTC and IIAI), Kai Zheng (UESTC), and Panos Kalnis (KAUST)</i>	
Time-Dependent Hop Labeling on Road Network	902
<i>Lei Li (The University of Queensland, Australia), Sibao Wang (The Chinese University of Hong Kong), and Xiaofang Zhou (The University of Queensland, Australia)</i>	
Weight-Constrained Route Planning Over Time-Dependent Graphs	914
<i>Ye Yuan (Beijing Institute of Technology), Xiang Lian (Kent State University), Guoren Wang (Beijing Institute of Technology), Lei Chen (Hong Kong University of Science and Technology), Yuliang Ma (Northeastern University, China), and Yishu Wang (Northeastern University, China)</i>	

Research (16) - Graphs 6: Data Science

Skyline Queries Constrained by Multi-cost Transportation Networks	926
<i>Qixu Gong (New Mexico State University), Huiping Cao (New Mexico State University), and Parth Nagarkar (New Mexico State University)</i>	
Online Social Media Recommendation Over Streams	938
<i>Xiangmin Zhou (RMIT University, Melbourne, Australia), Dong Qin (RMIT University, Melbourne, Australia), Xiaolu Lu (RMIT University, Melbourne, Australia), Lei Chen (Hong Kong University of Science and Technology, Hong Kong, China), and Yanchun Zhang (Victoria University, Melbourne, Australia; CIAT, Guangzhou University, China)</i>	
Canonicalization of Open Knowledge Bases with Side Information from the Source Text	950
<i>Xueling Lin (The Hong Kong University of Science and Technology) and Lei Chen (The Hong Kong University of Science and Technology)</i>	
Walking with Perception: Efficient Random Walk Sampling via Common Neighbor Awareness	962
<i>Yongkun Li (University of Science and Technology of China), Zhiyong Wu (University of Science and Technology of China), Shuai Lin (University of Science and Technology of China), Hong Xie (Chongqing University), Min Lv (University of Science and Technology of China), Yinlong Xu (University of Science and Technology of China), and John C.S. Lui (The Chinese University of Hong Kong)</i>	

Research (17) - Search and Information Extraction

SimMeme: A Search Engine for Internet Memes	974
<i>Tova Milo (Tel Aviv Univesity), Amit Somech (Tel Aviv Univesity), and Brit Youngmann (Tel Aviv Univesity)</i>	
A Hierarchical Framework for Top-k Location-Aware Error-Tolerant Keyword Search	986
<i>Junye Yang (Tsinghua University), Yong Zhang (Tsinghua University), Xiaofang Zhou (The University of Queensland), Jin Wang (University of California), Huiqi Hu (East China Normal University), and Chunxiao Xing (Tsinghua University)</i>	
2ED: An Efficient Entity Extraction Algorithm Using Two-Level Edit-Distance	998
<i>Zeyi Wen (National University of Singapore), Dong Deng (Rutgers University & Inception Institute of Artificial Intelligence), Rui Zhang (The University of Melbourne), and Ramamohanarao Kotagiri (The University of Melbourne)</i>	
Bridging Quantities in Tables and Text	1010
<i>Yusra Ibrahim (Max Planck Institute for Informatics), Mirek Riedewald (Northeastern University), Gerhard Weikum (Max Planck Institute for Informatics), and Demetrios Zeinalipour-Yazti (University of Cyprus)</i>	

Research (18) - Temporal and Spatial Data 2

An Efficient Insertion Operator in Dynamic Ridesharing Services	1022
<i>Yi Xu (Beihang University, China), Yongxin Tong (Beihang University, China), Yexuan Shi (Beihang University, China), Qian Tao (Beihang University, China), Ke Xu (Beihang University, China), and Wei Li (Beihang University, China)</i>	
Auction-Based Order Dispatch and Pricing in Ridesharing	1034
<i>Libin Zheng (The Hong Kong University of Science and Technology), Peng Cheng (The Hong Kong University of Science and Technology), and Lei Chen (The Hong Kong University of Science and Technology)</i>	
When Geo-Text Meets Security: Privacy-Preserving Boolean Spatial Keyword Queries	1046
<i>Ningning Cui (Northeastern University), Jianxin Li (Deakin University), Xiaochun Yang (Northeastern University), Bin Wang (Northeastern University), Mark Reynolds (University of Western Australia), and Yong Xiang (Deakin University)</i>	
Moving Object Linking Based on Historical Trace	1058
<i>Fengmei Jin (Renmin University of China; The University of Queensland), Wen Hua (The University of Queensland), Jiajie Xu (Soochow University; Neusoft Corporation), and Xiaofang Zhou (The University of Queensland; Soochow University)</i>	
ImageProof: Enabling Authentication for Large-Scale Image Retrieval	1070
<i>Shangwei Guo (Hong Kong Baptist University), Jianliang Xu (Hong Kong Baptist University), Ce Zhang (Hong Kong Baptist University), Cheng Xu (Hong Kong Baptist University), and Tao Xiang (Chongqing University)</i>	

Research (19) - Graphs 7: Mining 3

Time Constrained Continuous Subgraph Search Over Streaming Graphs	1082
<i>Youhuan Li (Center for Data Science, Peking University, China), Lei Zou (Institute of Computer Science and Technology, Peking University, China), M. Tamer Özsu (David R. Cheriton School of Computer Science), and Dongyan Zhao (Institute of Computer Science and Technology, Peking University, China)</i>	
Utilizing Dynamic Properties of Sharing Bits and Registers to Estimate User Cardinalities Over Time	1094
<i>Pinghui Wang (MOE Key Laboratory for Intelligent Networks and Network Security, Xi'an Jiaotong University, China; Shenzhen Research Institute of Xi'an Jiaotong University, Shenzhen, China), Peng Jia (MOE Key Laboratory for Intelligent Networks and Network Security, Xi'an Jiaotong University, China), Xiangliang Zhang (King Abdullah University of Science and Technology, Thuwal, SA), Jing Tao (MOE Key Laboratory for Intelligent Networks and Network Security, Xi'an Jiaotong University, China; Shenzhen Research Institute of Xi'an Jiaotong University, Shenzhen, China; Zhejiang Research Institute of Xi'an Jiaotong University, Hangzhou, China), Xiaohong Guan (Shenzhen Research Institute of Xi'an Jiaotong University, Shenzhen, China; MOE Key Laboratory for Intelligent Networks and Network Security, Xi'an Jiaotong University, China; Department of Automation and NLIST Lab, Tsinghua University, Beijing, China), and Don Towsley (School of Computer Science, University of Massachusetts Amherst, MA, USA)</i>	

Tracking Influential Nodes in Time-Decaying Dynamic Interaction Networks	1106
<i>Junzhou Zhao (King Abdullah University of Science and Technology), Shuo Shang (Inception Institute of Artificial Intelligence), Pinghui Wang (Xi'an Jiaotong University), John C.S. Lui (The Chinese University of Hong Kong), and Xiangliang Zhang (King Abdullah University of Science and Technology)</i>	
Fast and Accurate Graph Stream Summarization	1118
<i>Xiangyang Gou (Peking University), Lei Zou (Peking University), Chenxingyu Zhao (Peking University), and Tong Yang (Peking University)</i>	
Mining Periodic Cliques in Temporal Networks	1130
<i>Hongchao Qin (Northeastern University), Rong-Hua Li (Beijing Institute of Technology), Guoren Wang (Beijing Institute of Technology), Lu Qin (University of Technology Sydney), Yurong Cheng (Beijing Institute of Technology), and Ye Yuan (Northeastern University)</i>	

Research (20) - Data Integration, Probabilistic Databases

Coloring Embedder: A Memory Efficient Data Structure for Answering Multi-set Query	1142
<i>Yang Tong (Peking University), Dongsheng Yang (Peking University), Jie Jiang (Peking University), Siang Gao (Peking University), Bin Cui (Peking University), Lei Shi (Beihang University), and Xiaoming Li (Peking University)</i>	
Mining Order-Preserving Submatrices Under Data Uncertainty: A Possible-World Approach	1154
<i>Ji Cheng (The Hong Kong University of Science and Technology), Da Yan (The University of Alabama at Birmingham), Xiaotian Hao (The Hong Kong University of Science and Technology), and Wilfred Ng (The Hong Kong University of Science and Technology)</i>	
Multi-Dimensional Genomic Data Management for Region-Preserving Operations	1166
<i>Olha Horlova (Politecnico di Milano), Abdulrahman Kaitoua (DFKI & TU-Berlin), Volker Markl (DFKI & TU-Berlin), and Stefano Ceri (Politecnico di Milano)</i>	
Improved Algorithms for Maximal Clique Search in Uncertain Networks	1178
<i>Rong-Hua Li (Beijing Institute of Technology; National Engineering Laboratory for Big Data System Computing Technology), Qiangqiang Dai (Beijing Institute of Technology), Guoren Wang (Beijing Institute of Technology), Zhong Ming (Shenzhen University; National Engineering Laboratory for Big Data System Computing Technology), Lu Qin (University of Technology, Sydney), and Jeffrey Xu Yu (The Chinese University of Hong Kong)</i>	
Lazo: A Cardinality-Based Method for Coupled Estimation of Jaccard Similarity and Containment	1190
<i>Raul Castro Fernandez (MIT), Jisoo Min (MIT), Demitri Nava (MIT), and Samuel Madden (MIT)</i>	

Research (21) - Indexing

TARDIS: Distributed Indexing Framework for Big Time Series Data	1202
<i>Liang Zhang (Worcester Polytechnic Institute), Noura Alghamdi (Worcester Polytechnic Institute), Mohamed Y. Eltabakh (Worcester Polytechnic Institute), and Elke A. Rundensteiner (Worcester Polytechnic Institute)</i>	
Mostly Order Preserving Dictionaries	1214
<i>Chunwei Liu (University of Chicago), McKade Umbenhowe (University of Wyoming), Hao Jiang (University of Chicago), Pranav Subramaniam (University of Chicago), Jihong Ma (Alibaba Group), and Aaron J. Elmore (University of Chicago)</i>	
Multi-copy Cuckoo Hashing	1226
<i>Dagang Li (Peking University Shenzhen Graduate School), Rong Du (Peking University), Ziheng Liu (Peking University), Tong Yang (Peking University), and Bin Cui (Peking University)</i>	
Efficient Scalable Multi-attribute Index Selection Using Recursive Strategies	1238
<i>Rainer Schlosser (Hasso Plattner Institute), Jan Kossmann (Hasso Plattner Institute), and Martin Boissier (Hasso Plattner Institute)</i>	
To Index or Not to Index: Optimizing Exact Maximum Inner Product Search	1250
<i>Firas Abuzaied (Stanford DAWN Project), Geet Sethi (Stanford DAWN Project), Peter Bailis (Stanford DAWN Project), and Matei Zaharia (Stanford DAWN Project)</i>	

Research (22) - Road Networks

Distributed In-memory Trajectory Similarity Search and Join on Road Network	1262
<i>Haitao Yuan (Tsinghua University) and Guoliang Li (Tsinghua University)</i>	
Stochastic Weight Completion for Road Networks Using Graph Convolutional Networks	1274
<i>Jilin Hu (Aalborg University), Chenjuan Guo (Aalborg University), Bin Yang (Aalborg University), and Christian S. Jensen (Aalborg University)</i>	
Identifying the Most Interactive Object in Spatial Databases	1286
<i>Daichi Amagata (Osaka University) and Takahiro Hara (Osaka University)</i>	
An Efficient Framework for Correctness-Aware kNN Queries on Road Networks	1298
<i>Dan He (University of Queensland), Sibow Wang (the Chinese University of Hong Kong), Xiaofang Zhou (University of Queensland), and Reynold Cheng (the University of Hong Kong)</i>	
MPR — A Partitioning-Replication Framework for Multi-Processing kNN Search on Road Networks	1310
<i>Siqiang Luo (The University of Hong Kong), Ben Kao (The University of Hong Kong), Xiaowei Wu (University of Vienna), and Reynold Cheng (The University of Hong Kong)</i>	

Research (23) - Learning, Temporal and Spatial Data

AppUsage2Vec: Modeling Smartphone App Usage for Prediction	1322
<i>Sha Zhao (Zhejiang University), Zhiling Luo (Zhejiang University), Ziwen Jiang (Zhejiang University), Hatyan Wang (Zhejiang University), Feng Xu (Zhejiang University), Shijian Li (Zhejiang University), Jianwei Yin (Zhejiang University), and Gang Pan (Zhejiang University)</i>	
iFair: Learning Individually Fair Data Representations for Algorithmic Decision Making	1334
<i>Preethi Lahoti (Max Planck Institute for Informatics), Krishna P. Gummadi (Max Planck Institute for Software Systems), and Gerhard Weikum (Max Planck Institute for Informatics)</i>	
DBSCAN-MS: Distributed Density-Based Clustering in Metric Spaces	1346
<i>Keyu Yang (Zhejiang University), Yunjun Gao (Zhejiang University), Rui Ma (Zhejiang University), Lu Chen (Aalborg University), Sai Wu (Zhejiang University), and Gang Chen (Zhejiang University)</i>	
Computing Trajectory Similarity in Linear Time: A Generic Seed-Guided Neural Metric Learning Approach	1358
<i>Di Yao (Institute of Computing Technology, Chinese Academy of Sciences), Gao Cong (Nanyang Technological University, Singapore), Chao Zhang (University of Illinois at Urbana-Champaign), and Jingping Bi (Institute of Computing Technology, Chinese Academy of Sciences)</i>	
Bursty Event Detection Throughout Histories	1370
<i>Debjyoti Paul (University of Utah), Yanqing Peng (University of Utah), and Feifei Li (University of Utah)</i>	

Research (24) - Knowledge Discovery

RUM: Network Representation Learning Using Motifs	1382
<i>Yanlei Yu (Renmin University of China), Zhiwu Lu (Renmin University of China), Jiajun Liu (Renmin University of China), Guoping Zhao (Renmin University of China), and Ji-rong Wen (Renmin University of China)</i>	
Finding Significant Items in Data Streams	1394
<i>Tong Yang (Peking University, China), Haowei Zhang (Peking University, China), Dongsheng Yang (Peking University, China), Yucheng Huang (Peking University, China), and Xiaoming Li (Peking University, China)</i>	
Knowledge-Aware Deep Dual Networks for Text-Based Mortality Prediction	1406
<i>Ning Liu (Tsinghua University), Pan Lu (Tsinghua University), Wei Zhang (East China Normal University), and Jianyong Wang (Tsinghua University)</i>	
Robust High Dimensional Stream Classification with Novel Class Detection	1418
<i>Zhuoyi Wang (University of Texas at Dallas), Zelun Kong (University of Texas at Dallas), Swarup Changra (University of Texas at Dallas), Hemeng Tao (University of Texas at Dallas), and Latifur Khan (University of Texas at Dallas)</i>	
Towards the Completion of a Domain-Specific Knowledge Base with Emerging Query Terms	1430
<i>Sihang Jiang (Fudan University), Jiaqing Liang (Fudan University), Yanghua Xiao (Fudan University), Haihong Tang (Alibaba Group), Haikuan Huang (Alibaba Group), and Jun Tan (Alibaba Group)</i>	

Research (25) - Crowdsourcing and Learning

Cooperation-Aware Task Assignment in Spatial Crowdsourcing	1442
<i>Peng Cheng (HKUST), Lei Chen (HKUST), and Jieping Ye (AI Labs, DiDi Chuxing)</i>	
Minimizing Maximum Delay of Task Assignment in Spatial Crowdsourcing	1454
<i>Zhao Chen (HKUST), Peng Cheng (HKUST), Yuxiang Zeng (HKUST), and Lei Chen (HKUST)</i>	
Physical Representation-Based Predicate Optimization for a Visual Analytics Database	1466
<i>Michael R. Anderson (University of Michigan), Michael Cafarella (University of Michigan), German Ros (Intel Labs), and Thomas F. Wensich (University of Michigan)</i>	
Adaptive Dynamic Bipartite Graph Matching: A Reinforcement Learning Approach	1478
<i>Yansheng Wang (Beihang University), Yongxin Tong (Beihang University), Cheng Long (Nanyang Technological University), Pan Xu (University of Maryland, College Park), Ke Xu (Beihang University), and Weifeng Lv (Beihang University)</i>	

Research (26) - Data Mining

Scalable Frequent Sequence Mining with Flexible Subsequence Constraints	1490
<i>Alexander Renz-Wieland (Technische Universität Berlin), Matthias Bertsch (Universität Mannheim), and Rainer Gemulla (Universität Mannheim)</i>	
Adaptive Influence Blocking: Minimizing the Negative Spread by Observation-Based Policies	1502
<i>Qihao Shi (Zhejiang University), Can Wang (Zhejiang University), Deshi Ye (Zhejiang University), Jiawei Chen (Zhejiang University), Yan Feng (Zhejiang University), and Chun Chen (Zhejiang University)</i>	
Fraction-Score: A New Support Measure for Co-location Pattern Mining	1514
<i>Harry Kai-Ho Chan (The Hong Kong University of Science and Technology), Cheng Long (Nanyang Technological University), Da Yan (The University of Alabama at Birmingham), and Raymond Chi-Wing Wong (The Hong Kong University of Science and Technology)</i>	
Discovery and Ranking of Functional Dependencies	1526
<i>Ziheng Wei (The University of Auckland) and Sebastian Link (The University of Auckland)</i>	
Adaptive Deep Reuse: Accelerating CNN Training on the Fly	1538
<i>Lin Ning (North Carolina State University), Hui Guan (North Carolina State University), and Xipeng Shen (North Carolina State University)</i>	

Research Short Papers

Slice Finder: Automated Data Slicing for Model Validation	1550
<i>Yeounoh Chung (Brown University), Tim Kraska (MIT), Neoklis Polyzotis (Google Research), Ki Hyun Tae (KAIST), and Steven Euijong Whang (KAIST)</i>	

Neural Multi-task Recommendation from Multi-behavior Data	1554
<i>Chen Gao (Tsinghua University), Xiangnan He (University of Science and Technology of China), Dahua Gan (Tsinghua University), Xiangning Chen (Tsinghua University), Fuli Feng (National University of Singapore), Yong Li (Tsinghua University), Tat-Seng Chua (National University of Singapore), and Depeng Jin (Tsinghua University)</i>	
AUC-MF: Point of Interest Recommendation with AUC Maximization	1558
<i>Peng Han (King Abdullah University of Science and Technology), Shuo Shang (Inception Institute of Artificial Intelligence), Aixin Sun (Nanyang Technological University), Peilin Zhao (Tencent AI Lab), Kai Zheng (University of Electronic Science and Technology of China), and Panos Kalnis (King Abdullah University of Science and Technology)</i>	
Efficient Batch One-Hop Personalized PageRanks	1562
<i>Siqiang Luo (The University of Hong Kong), Xiaokui Xiao (National University of Singapore), Wenqing Lin (Tencent), and Ben Kao (The University of Hong Kong)</i>	
BB-Tree: A Main-Memory Index Structure for Multidimensional Range Queries	1566
<i>Stefan Sprenger (Humboldt-University Berlin), Patrick Schäfer (Humboldt-University Berlin), and Ulf Leser (Humboldt-University Berlin)</i>	
Explaining Queries Over Web Tables to Non-experts	1570
<i>Jonathan Berant (Tel Aviv University), Daniel Deutch (Tel Aviv University), Amir Globerson (Tel Aviv University), Tova Milo (Tel Aviv University), and Tomer Wolfson (Tel Aviv University)</i>	
Nonlinear Models Over Normalized Data	1574
<i>Zhaoyue Cheng (University of Toronto) and Nick Koudas (University of Toronto)</i>	
Continuous Search on Dynamic Spatial Keyword Objects	1578
<i>Yuyang Dong (University of Tsukuba), Hanxiong Chen (University of Tsukuba), and Hiroyuki Kitagawa (University of Tsukuba)</i>	
Top-K Frequent Term Queries on Streaming Data	1582
<i>Sara Farazi (University of Alberta) and Davood Rafiei (University of Alberta)</i>	
Parallel and Distributed Processing of Reverse Top-k Queries	1586
<i>Panagiotis Nikitopoulos (University of Piraeus), Georgios A. Sfyris (University of Piraeus), Akrivi Vlachou (University of Piraeus), Christos Doulkeridis (University of Piraeus), and Orestis Telelis (University of Piraeus)</i>	
Distributed Discovery of Functional Dependencies	1590
<i>Hemant Saxena (University of Waterloo), Lukasz Golab (University of Waterloo), and Ihab F. Ilyas (University of Waterloo)</i>	
AID: An Adaptive Image Data Index for Interactive Multilevel Visualization	1594
<i>Saheli Ghosh (University of California, Riverside), Ahmed Eldawy (University of California, Riverside), and Shipra Jais (University of California, Riverside)</i>	

Collecting Preference Rankings Under Local Differential Privacy	1598
<i>Jianyu Yang (Beijing University of Posts and Telecommunications), Xiang Cheng (Beijing University of Posts and Telecommunications), Sen Su (Beijing University of Posts and Telecommunications), Rui Chen (Samsung Research America, Mountain View, USA), Qiyu Ren (Beijing University of Posts and Telecommunications), and Yuhan Liu (Beijing University of Posts and Telecommunications)</i>	
Muses: Distributed Data Migration System for Polystores	1602
<i>Abdulrahman Kaitoua (DFKI), Tilmann Rabl (TU-Berlin), Asterios Katsifodimos (Delft University of Technology), and Volker Markl (TU-Berlin)</i>	
PriSTE: From Location Privacy to Spatiotemporal Event Privacy	1606
<i>Yang Cao (Kyoto University), Yonghui Xiao (Google Inc.), Li Xiong (Emory University), and Liquan Bai (Emory University)</i>	
Continuous Range Queries Over Multi-attribute Trajectories	1610
<i>jianqiu xu (Nanjing University of Aeronautics and Astronautics, China), Zhifeng Bao (RMIT University, Australia), and Hua Lu (Aalborg University, Denmark)</i>	
Insecurity and Hardness of Nearest Neighbor Queries Over Encrypted Data	1614
<i>Rui Li (Dongguan University of Technology), Alex X. Liu (Dongguan University of Technology), Ying Liu (Hunan University), Huanle Xu (Dongguan University of Technology), and Huaqiang Yuan (Dongguan University of Technology)</i>	
Modeling Multidimensional User Preferences for Collaborative Filtering	1618
<i>Farhan Khawar (The Hong Kong University of Science and Technology) and Nevin L. Zhang (The Hong Kong University of Science and Technology)</i>	
A Queueing-Theoretic Framework for Vehicle Dispatching in Dynamic Car-Hailing	1622
<i>Peng Cheng (HKUST), Chao Feng (AI Labs, DiDi Chuxing), Lei Chen (HKUST), and Zheng Wang (AI Labs, DiDi Chuxing)</i>	
Maximizing the Utility in Location-Based Mobile Advertising	1626
<i>Peng Cheng (HKUST), Xiang Lian (Kent State University), Lei Chen (HKUST), and Siyuan Liu (Pennsylvania State University)</i>	
Automated Grading of SQL Queries	1630
<i>Bikash Chandra (IIT Bombay), Ananyo Banerjee (IIT Bombay), Udbhas Hazra (IIT Bombay), Mathew Joseph (IIT Bombay), and S. Sudarshan (IIT Bombay)</i>	
Parameter Discovery in Unsupervised Clustering	1634
<i>Valentin Clement (Imperial College London) and Thomas Heinis (Imperial College London)</i>	
Interaction-Aware Arrangement for Event-Based Social Networks	1638
<i>Feifei Kou (Beijing University of Posts and Telecommunications), Zimu Zhou (ETH Zurich), Hao Cheng (Beihang University), Junping Du (Beijing University of Posts and Telecommunications), Yexuan Shi (Beihang University), and Pan Xu (University of Maryland)</i>	
Optimizing Cross-Platform Data Movement	1642
<i>Sebastian Kruse (HPI), Zoi Kaoudi (QCRI), Jorge-Arnulfo Quijane-Ruiz (QCRI), Sanjay Chawla (QCRI), Felix Naumann (QCRI), and Bertty Contreras-Rojas (QCRI)</i>	

Highly Efficient Pattern Mining Based on Transaction Decomposition	1646
<i>Yousef Djenouri (NTNU, Trondheim, Norway), Jerry Chun-Wei Lin (HVL, Bergen, Norway), Kjetil Nørnvåg (NTNU, Trondheim, Norway), and Heri Ramampiaro (HVL, Bergen, Norway)</i>	
Procrastination-Aware Scheduling: A Bipartite Graph Perspective	1650
<i>Libin Wang (Beihang University), Yongxin Tong (Beihang University), Chunming Hu (Beihang University), Lei Chen (The Hong Kong University of Science and Technology), and Yiming Li (Beihang University)</i>	
Hankel Matrix Factorization for Tagged Time Series to Recover Missing Values During Blackouts	1654
<i>Simeng Wu (State Key Laboratory for Novel Software Technology, Nanjing University), Liang Wang (State Key Laboratory for Novel Software Technology, Nanjing University), Tianheng Wu (State Key Laboratory for Novel Software Technology, Nanjing University), Xianping Tao (State Key Laboratory for Novel Software Technology, Nanjing University), and Jian Lu (State Key Laboratory for Novel Software Technology, Nanjing University)</i>	
PerRD: A System for Personalized Route Description	1658
<i>Han Su (Zhejiang Gongshang University), Guanglin Cong (University of Electronic Science and Technology of China), Wei Chen (University of Electronic Science and Technology of China), Qinyuan Su (University of Electronic Science and Technology of China), Bolong Zheng (Huazhong University of Science and Technology), and Kai Zheng (University of Electronic Science and Technology of China)</i>	
Scalable Metric Similarity Join Using MapReduce	1662
<i>Jiacheng Wu (Tsinghua University), Yong Zhang (Tsinghua University), Jin Wang (University of California, Los Angeles), Chunbin Lin (Amazon AWS), Yingjia Fu (University of California, San Diego), and Chunxiao Xing (Tsinghua University)</i>	
An Indexing Framework for Efficient Visual Exploratory Subgraph Search in Graph Databases	1666
<i>Chaohui Wang (Nanyang Technological University), Miao Xie (Alibaba Group), Sourav S. Bhowmick (Nanyang Technological University), Byron Choi (Hong Kong Baptist University), Xiaokui Xiao (National University of Singapore), and Shuigeng Zhou (Fudan University)</i>	
I-LSH: I/O Efficient c-Approximate Nearest Neighbor Search in High-Dimensional Space	1670
<i>Wanqi Liu (University of Technology Sydney), Hanchen Wang (University of Technology Sydney), Ying Zhang (University of Technology Sydney), Wei Wang (university of new south wales), and Lu Qin (University of Technology Sydney)</i>	
T-Sample: A Dual Reservoir-Based Sampling Method for Characterizing Large Graph Streams	1674
<i>Lingling Zhang (Huazhong University of Science and Technology), Hong Jiang (Department of Computer Science and Engineering, University of Texas at Arlington, USA), Fang Wang (Huazhong University of Science and Technology), Dan Feng (Huazhong University of Science and Technology), and Yanwen Xie (Huazhong University of Science and Technology)</i>	

Real Time Principal Component Analysis	1678
<i>Ranak Roy Chowdhury (Bangladesh University of Engineering and Technology (BUET)), Muhammad Abdullah Adnan (Bangladesh University of Engineering and Technology (BUET)), and Rajesh K. Gupta (University of California San Diego)</i>	
A Fast Sketch Method for Mining User Similarities Over Fully Dynamic Graph Streams	1682
<i>Peng Jia (Xi'an Jiaotong University), Pinghui Wang (Xi'an Jiaotong University), Jing Tao (Xi'an Jiaotong University), and Xiaohong Guan (Xi'an Jiaotong University)</i>	
A Collaborative Framework for Similarity Enforcement in Synthetic Scaling of Relational Datasets	1686
<i>Jiangwei W. Zhang (National University of Singapore) and Yong Chiang Tay (National University of Singapore)</i>	
Meta Diagram Based Active Social Networks Alignment	1690
<i>Yuxiang Ren (IFM lab, Florida State University), Charu C. Aggarwal (IBM Research AI), and Jiawei Zhang (IFM lab, Florida State University)</i>	
Entity Integrity, Referential Integrity, and Query Optimization with Embedded Uniqueness Constraints.....	1694
<i>Ziheng Wei (The University of Auckland), Uwe Leck (The University of Flensburg), and Sebastian Link (The University of Auckland)</i>	
Efficient Pattern Mining Based Cryptanalysis for Privacy-Preserving Record Linkage	1698
<i>Anushka Vidanage (The Australian National University, Canberra, Australia), Thilina Ranbaduge (The Australian National University, Canberra, Australia), Peter Christen (The Australian National University, Canberra, Australia), and Rainer Schnell (University Duisburg-Essen, Germany)</i>	
ECOQUG: An Effective Ensemble Community Scoring Function	1702
<i>Chunnan Wang (Harbin Institute of Technology), Hongzhi Wang (Harbin Institute of Technology), Chang Zhou (Harbin Institute of Technology), Jianzhong Li (Harbin Institute of Technology), and Hong Gao (Harbin Institute of Technology)</i>	
CN-Probase: A Data-Driven Approach for Large-Scale Chinese Taxonomy Construction	1706
<i>Jindong Chen (Fudan University), Ao Wang (Fudan University), Jiangjie Chen (Fudan University), Yanghua Xiao (Fudan University), Zhendong Chu (Fudan University), Jingping Liu (Fudan University), Jiaqing Liang (Fudan University), and Wei Wang (Fudan University)</i>	
Understanding Data Science Lifecycle Provenance via Graph Segmentation and Summarization	1710
<i>Hui Miao (University of Maryland) and Amol Deshpande (University of Maryland)</i>	
Efficient Partitioning and Query Processing of Spatio-Temporal Graphs with Trillion Edges	1714
<i>Mengsu Ding (SKL Computer Architecture, ICT, CAS, University of Chinese Academy of Sciences) and Shimin Chen (SKL Computer Architecture, ICT, CAS, University of Chinese Academy of Sciences)</i>	
Top-K Entity Resolution with Adaptive Locality-Sensitive Hashing	1718
<i>Vasilis Verroios (Stanford University) and Hector Garcia-Molina</i>	
Finding Average Regret Ratio Minimizing Set in Database	1722
<i>Sepanta Zeighami (Hong Kong University of Science and Technology) and Raymond Chi-Wing Wong (Hong Kong University of Science and Technology)</i>	

HyMJ: A Hybrid Structure-Aware Approach to Distributed Multi-way Join Query	1726
<i>Guanghai Zhu (Nanjing University), Xiaoqi Wu (Nanjing University), Liangliang Yin (Nanjing University), Haogang Wang (Nanjing University), Rong Gu (Nanjing University), Chunfeng Yuan (Nanjing University), and Yihua Huang (Nanjing University)</i>	
Accelerate MaxBRkNN Search by kNN Estimation	1730
<i>Xuefeng Chen (University of New South Wales, Australia), Xin Cao (University of New South Wales, Australia), Zhiqiang Xu (Baidu, China), Ying Zhang (University of Technology Sydney, Australia), Shuo Shang (Inception Institute of Artificial Intelligence, United Arab Emirates), and Wenjie Zhang (University of New South Wales, Australia)</i>	
Efficient Bottom-Up Discovery of Multi-scale Time Series Correlations Using Mutual Information	1734
<i>Nguyen Ho (Aalborg University), Torben Bach Pedersen (Aalborg University), Mai Vu (Tufts University), Van Long Ho (Aalborg University), and Christophe A.N. Biscio (Aalborg University)</i>	
Fingerprinting Big Data: The Case of KNN Graph Construction	1738
<i>Rachid Guerraoui (EPFL), Anne-Marie Kermarrec (EPFL, Mediego), Olivier Ruas (Inria, Univ Rennes, CNRS, IRISA), and François Taïani (Univ Rennes, Inria, CNRS, IRISA)</i>	
Outer and Anti Joins in Temporal-Probabilistic Databases	1742
<i>Katerina Papaioannou (University of Zurich), Martin Theobald (University of Luxembourg), and Michael Boehlen (University of Zurich)</i>	
Workload-Driven Fragment Allocation for Partially Replicated Databases Using Linear Programming	1746
<i>Stefan Halfpap (Hasso Plattner Institute, University of Potsdam, Germany) and Rainer Schlosser (Hasso Plattner Institute, University of Potsdam, Germany)</i>	
OSMAC: Optimizing Subgraph Matching Algorithms with Community Structure	1750
<i>Yunkai Lou (Tsinghua University) and Chaokun Wang (Tsinghua University)</i>	
Workload-Aware Subgraph Query Caching and Processing in Large Graphs	1754
<i>Yongjiang Liang (Florida State University) and Peixiang Zhao (Florida State University)</i>	
How I Learned to Stop Worrying and Love Re-optimization	1758
<i>Matthew Perron (MIT CSAIL), Zeyuan Shang (MIT CSAIL), Tim Kraska (MIT CSAIL), and Michael Stonebraker (MIT CSAIL)</i>	
CRA: Enabling Data-Intensive Applications in Containerized Environments	1762
<i>Ibrahim Sabek (University of Minnesota, USA), Badrish Chandramouli (Microsoft Research, USA), and Umar Farooq Minhas (Microsoft Research, USA)</i>	

Industry Track

Industry (1) - Data Systems Optimizations

Scalable Similarity Joins of Tokenized Strings	1766
<i>Ahmed Metwally (LinkedIn) and Chun-Heng Huang (Google)</i>	

MLlib*: Fast Training of GLMs Using Spark MLlib	1778
<i>Zhipeng Zhang (Peking University), Jiawei Jiang (Tencent Inc.), Wentao Wu (Microsoft Research, Redmond, USA), Ce Zhang (ETH Zurich, Switzerland), Lele Yu (Tencent Inc.), and Bin Cui (Peking University)</i>	
DirectLoad: A Fast Web-Scale Index System Across Large Regional Centers	1790
<i>An Qin (Baidu Inc.), Mengbai Xiao (The Ohio State University), Jin Ma (Baidu Inc.), Dai Tan (Baidu Inc.), Rubao Lee (United Parallel Computing Corporation), and Xiaodong Zhang (The Ohio State University)</i>	
Presto: SQL on Everything	1802
<i>Raghav Sethi (Facebook, Inc.), Martin Traverso (Facebook, Inc.), Dain Sundstrom (Facebook, Inc.), David Phillips (Facebook, Inc.), Wenlei Xie (Facebook, Inc.), Yutian Sun (Facebook, Inc.), Nezh Yegitbasi (Facebook, Inc.), Haozhun Jin (Facebook, Inc.), Eric Hwang (Facebook, Inc.), Nileema Shingte (Facebook, Inc.), and Christopher Berner (Facebook, Inc.)</i>	
Improving RDF Query Performance Using In-memory Virtual Columns in Oracle Database	1814
<i>Eugene Inseok Chong (Oracle), Matthew Perry (Oracle), and Souripriya Das (Oracle)</i>	

Industry (2) - Data Analytics

SEBDB: Semantics Empowered Blockchain DataBase	1820
<i>Yanchao Zhu (East China Normal University), Zhao Zhang (East China Normal University), Cheqing Jin (East China Normal University), Aoying Zhou (East China Normal University), and Ying Yan (Ant Financial)</i>	
Large Scale Traffic Signal Network Optimization - A Paradigm Shift Driven by Big Data	1832
<i>Liang Yu (Alibaba Cloud), Jinqiang Yu (Alibaba Cloud), Maolei Zhang (Alibaba Cloud), Xin Zhang (Alibaba Cloud), Yuehu Liu (Alibaba Cloud), Hui Zhang (Alibaba Cloud), and Wanli Min (Alibaba Cloud)</i>	
Domain-Independent Automated Processing of Free-Form Text Data in Telecom	1841
<i>Rajarshi Bhowmik (Rutgers University) and Ahmet Akyamac (Bell Labs, Nokia)</i>	
DRIVEN: a Framework for Efficient Data Retrieval and Clustering in Vehicular Networks	1850
<i>Bastian Havers (Chalmers University), Romaric Duvignau (Chalmers University), Hannaneh Najdataei (Chalmers University), Vincenzo Gulisano (Chalmers University), Ashok Chaitanya Koppisetty (Volvo Cars), and Marina Papatriantafilou (Chalmers University)</i>	

Industry (3) - Driving Business Applications

Accurate Product Attribute Extraction on the Field	1862
<i>Martin Rezk (Rakuten Inc), Laura Alonso Alemany (Universidad Nacional de Cordoba), Lasguido Nio (Rakuten Inc), and Ted Zhang (Rakuten Inc)</i>	

CATS: Cross-Platform E-Commerce Fraud Detection	1874
<i>Haiqin Weng (Zhejiang University), Shouling Ji (Zhejiang University), Fuzheng Duan (Zhejiang University), Zhao Li (Alibaba Group), Jianhai Chen (Zhejiang University), Qinming He (Zhejiang University), and Ting Wang (Lehigh University)</i>	
Caladrius: A Performance Modelling Service for Distributed Stream Processing Systems	1886
<i>Faria Kalim (University of Illinois at Urbana-Champaign), Thomas Cooper (Newcastle University, UK), Huijun Wu (Twitter, Inc.), Yao Li (Twitter, Inc.), Ning Wang (Twitter, Inc.), Neng Lu (Twitter, Inc.), Maosong Fu (Twitter, Inc.), Xiaoyao Qian (Twitter, Inc.), Hao Luo (Twitter, Inc.), Da Cheng (Twitter, Inc.), Yaliang Wang (Twitter, Inc.), Fred Dai (Twitter, Inc.), Mainak Ghosh (Twitter, Inc.), and Beinan Wang (Twitter, Inc.)</i>	
FAIR: Fraud Aware Impression Regulation System in Large-Scale Real-Time E-Commerce Search Platform .	1898
<i>Zhao Li (Alibaba Group), Junshuai Song (Peking University), Shichang Hu (Alibaba Group), Shasha Ruan (Alibaba Group), Long Zhang (Alibaba Group), Zehong Hu (Alibaba Group), and Jun Gao (Peking University)</i>	
Micro-Browsing Models for Search Snippets	1904
<i>Muhammad Asiful Islam (Google Inc), Ramakrishnan Srikant (Google Inc), and Sugato Basu (Google Inc)</i>	

Industry (4) - Learning from Data

Interpretable Multi-task Learning for Product Quality Prediction with Attention Mechanism	1910
<i>Cheng-Han Yeh (National Chiao Tung University, Hsinchu, Taiwan), Yao-Chung Fan (National Chung Hsing University, Taichung, Taiwan), and Wen-Chih Peng (National Chiao Tung University, Hsinchu, Taiwan)</i>	
Learning Effective Embeddings From Crowdsourced Labels: An Educational Case Study	1922
<i>Guowei Xu (TAL AI Lab), Wenbiao Ding (TAL AI Lab), Jiliang Tang (Michigan State University), Songfan Yang (TAL AI Lab), Gale Yan Huang (TAL AI Lab), and Zitao Liu (TAL AI Lab)</i>	
A Prescription Trend Analysis using Medical Insurance Claim Big Data	1928
<i>Kazutoshi Umemoto (The University of Tokyo), Kazuo Goda (The University of Tokyo), Naohiro Mitsutake (Institute for Health Economics and Policy), and Masaru Kitsuregawa (The University of Tokyo)</i>	
Differential Data Quality Verification on Partitioned Data	1940
<i>Sebastian Schelter (Amazon Research), Stefan Grafberger (Amazon Research), Philipp Schmidt (Amazon Research), Tammo Rukat (Amazon Research), Mario Kiessling (Amazon Research), Andrey Taptunov (Amazon Research), Felix Biessmann (Amazon Research), and Dustin Lange (Amazon Research)</i>	
Logan: A Distributed Online Log Parser	1946
<i>Amey Agrawal (Qubole India), Rohit Karlupia (Qubole India), and Rajat Gupta (Qubole India)</i>	

Demo Track

Demo Session 1

WebPut: A Web-Aided Data Imputation System for the General Type of Missing String Attribute Values	1952
<i>Shuangli Shan (Soochow University, China; Neusoft Corporation), Zhixu Li (Soochow University, China), Yang Li (Soochow University, China), Qiang Yang (King Abdullah University of Science and Technology, Saudi Arabia), Jia Zhu (South China Normal University), Mohamed Sharaf (The University of Queensland), and Xiaofang Zhou (The University of Queensland)</i>	
FGreat: Focused Graph Query Autocompletion	1956
<i>Nathan Ng (Hong Kong Baptist University), Peipei Yi (Lenovo Machine Intelligence Center), Zhiwei Zhang (Hong Kong Baptist University), Byron Choi (Hong Kong Baptist University), Sourav S. Bhowmick (Nanyang Technological University), and Jianliang Xu (Hong Kong Baptist University)</i>	
Aucher: Multi-modal Queries on Live Audio Streams in Real-Time	1960
<i>Zeyi Wen (National University of Singapore), Mingyu Liang (Shanghai Jiao Tong University), Bingsheng He (National University of Singapore), Zexin Xia (Shanghai Jiao Tong University), and Bo Li (Hong Kong University of Science and Technology)</i>	
SAC: A System for Big Data Lineage Tracking	1964
<i>Mingjie Tang (Hortonworks), Saisai Shao (Hortonworks), Weiqing Yang (Hortonworks), Yanbo Liang (Hortonworks), Yongyang Yu (Facebook), Bikas Saha (Hortonworks), and Dongjoon Hyun (Hortonworks)</i>	
A Gossip-Based System for Fast Approximate Score Computation in Multinomial Bayesian Networks	1968
<i>Arun Zachariah (University of Missouri-Kansas City), Praveen Rao (University of Missouri-Kansas City), Anas Katib (University of Missouri-Kansas City), Monica Senapati (University of Missouri-Kansas City), and Kobus Barnard (University of Arizona)</i>	
Faster, Higher, Stronger: Redesigning Spreadsheets for Scale	1972
<i>Mangesh Bendre (University of Illinois at Urbana-Champaign), Tana Wattanawaroon (University of Illinois at Urbana-Champaign), Sajjadur Rahman (University of Illinois at Urbana-Champaign), Kelly Mack (University of Illinois at Urbana-Champaign), Yuyang Liu (University of Illinois at Urbana-Champaign), Shichu Zhu (University of Illinois at Urbana-Champaign), Yu Lu (University of Illinois at Urbana-Champaign), Ping-Jing Yang (University of Illinois at Urbana-Champaign), Xinyan Zhou (University of Illinois at Urbana-Champaign), Kevin Chen-Chuan Chang (University of Illinois at Urbana-Champaign), Karrie Karahalios (University of Illinois at Urbana-Champaign), and Aditya Parameswaran (University of Illinois at Urbana-Champaign)</i>	
RecovDB: Accurate and Efficient Missing Blocks Recovery for Large Time Series	1976
<i>Ines Arous (University of Fribourg), Mourad Khayati (University of Fribourg), Philippe Cudré-Mauroux (University of Fribourg), Ying Zhang (MonetDB Solutions), Martin Kersten (MonetDB Solutions), and Svetlin Stalnikov (MonetDB Solutions)</i>	

AI Pro: Data Processing Framework for AI Models	1980
<i>Richie Frost (University of Utah), Debjyoti Paul (University of Utah), and Feifei Li (University of Utah)</i>	
IVLG: Interactive Visualization of Large Graphs	1984
<i>Maria Krommyda (NTUA), Verena Kantere (University of Ottawa), and Yannis Vassiliou (NTUA)</i>	
Just in Time: Personal Temporal Insights for Altering Model Decisions	1988
<i>Naama Boer (Tel Aviv University), Daniel Deutch (Tel Aviv University), Nave Frost (Tel Aviv University), and Tova Milo (Tel Aviv University)</i>	
GeoSparkViz in Action: A Data System with Built-in Support for Geospatial Visualization	1992
<i>Jia Yu (Arizona State University), Anique Tahir (Arizona State University), and Mohamed Sarwat (Arizona State University)</i>	
Hybrid.Poly: A Consolidated Interactive Analytical Polystore System	1996
<i>Maksim Podkorytov (Florida State University) and Michael Gubanov (Florida State University)</i>	

Demo Session 2

EXPLAINER: Entity Resolution Explanations	2000
<i>Amr Ebaid (Purdue University), Saravanan Thirumuruganathan (Qatar Computing Research Institute, HBKU), Walid G. Aref (Purdue University), Ahmed Elmagarmid (Qatar Computing Research Institute, HBKU), and Mourad Ouzzani (Qatar Computing Research Institute, HBKU)</i>	
CEP-Wizard: Automatic Deployment of Distributed Complex Event Processing	2004
<i>Yooju Shin (KAIST), Susik Yoon (KAIST), Patara Trirat (KAIST), and Jae-Gil Lee (KAIST)</i>	
A Comparison of Allocation Algorithms for Partially Replicated Databases	2008
<i>Stefan Halfpap (Hasso Plattner Institute, University of Potsdam, Germany) and Rainer Schlosser (Hasso Plattner Institute, University of Potsdam, Germany)</i>	
PePPER: Fine-Grained Personal Access Control via Peer Probing	2012
<i>Yael Amsterdamer (Bar-Ilan University) and Osnat Drien (Bar-Ilan University)</i>	
COBRA: Compression Via Abstraction of Provenance for Hypothetical Reasoning	2016
<i>Daniel Deutch (Tel Aviv University), Yuval Moskovitch (Tel Aviv University), and Noam Rinetzky (Tel Aviv University)</i>	
CogLearn: A Cognitive Graph-Oriented Online Learning System	2020
<i>Yang Pian (Beijing Normal University), Yu Lu (Beijing Normal University), Penghe Chen (Beijing Normal University), and Qinglong Duan (Beijing Normal University)</i>	
GRIT: Consistent Distributed Transactions Across Polyglot Microservices with Multiple Databases	2024
<i>Guogen Zhang (eBay Inc.), Kun Ren (eBay Inc.), Jung-Sang Ahn (eBay Inc.), and Sami Ben-Romdhane (eBay Inc.)</i>	
vABS: Towards Verifiable Attribute-Based Search Over Shared Cloud Data	2028
<i>Yang Ji (Hong Kong Baptist University), Cheng Xu (Hong Kong Baptist University), Jianliang Xu (Hong Kong Baptist University), and Haibo Hu (Hong Kong Polytechnic University)</i>	

An Environment-Aware Market Strategy for Data Allocation and Dynamic Migration in Cloud Database	2032
<i>Tengjiao Wang (Peking University), Binyang Li (University of International Relations), Wei Chen (Peking University), Yuxiao Zhang (Peking University), Ying Han (Peking University), Jinzhong Niu (City University of New York), and Kam-fai Wong (The Chinese University of Hong Kong)</i>	
Vaite: A Visualization-Assisted Interactive Big Urban Trajectory Data Exploration System	2036
<i>Chuang Yang (Southern University of Science and Technology; PCL Research Center of Networks and Communications, Peng Cheng Laboratory), Yilan Zhang (Sichuan University), Bo Tang (Southern University of Science and Technology; PCL Research Center of Networks and Communications, Peng Cheng Laboratory), and Min Zhu (Sichuan University)</i>	
SciDetector: Scientific Event Discovery by Tracking Variable Source Data Streaming	2040
<i>Zhiqiang Duan (Renmin University of China), Chen Yang (Renmin University of China), Xiaofeng Meng (Renmin University of China), Yongjie Du (Renmin University of China), Jiaming Qiu (Tsinghua University), Xiaobin Ma (Tsinghua University), Zhihui Du (Tsinghua University), Xukang Zhang (Taiyuan University of Technology), Baoning Niu (Taiyuan University of Technology), and Chao Wu (National Astronomical Observatories, Chinese Academy of Sciences)</i>	
Demonstrating Spindra: A Geographic Knowledge Graph Management System	2044
<i>Yuhan Sun (Arizona State University), Jia Yu (Arizona State University), and Mohamed Sarwat (Arizona State University)</i>	

Tutorials

Native Storage Techniques for Data Management	2048
<i>Iliia Petrov (Reutlingen University), Andreas Koch (Technische Universität Darmstadt), Sergey Hardock (Technische Universität Darmstadt), Tobias Vincon (Reutlingen University), and Christian Riegger (Reutlingen University)</i>	
Crowdsourcing Database Systems: Overview and Challenges	2052
<i>Chengliang Chai (Tsinghua University), Ju Fan (Renmin University), Guoliang Li (Tsinghua University), Jiannan Wang (Simon Fraser University), and Yudian Zheng (Twitter)</i>	
Telco Big Data Research and Open Problems	2056
<i>Constantinos Costa (University of Pittsburgh) and Demetrios Zeinalipour-Yazti (University of Cyprus)</i>	
Geospatial Data Management in Apache Spark: A Tutorial	2060
<i>Jia Yu (Arizona State University) and Mohamed Sarwat (Arizona State University)</i>	
Hierarchical Decomposition of Big Graphs	2064
<i>Ying Zhang (University of Technology Sydney), Lu Qin (University of Technology Sydney), Fan Zhang (University of New South Wales), and Wenjie Zhang (University of New South Wales)</i>	

Cohesive Subgraph Computation Over Large Sparse Graphs	2068
<i>Lijun Chang (The University of Sydney) and Lu Qin (University of Technology Sydney)</i>	
Robust Query Processing: Mission Possible	2072
<i>Jayant R. Haritsa (Indian Institute of Science)</i>	

PhD Symposium

Automated Documentation of End-to-End Experiments in Data Science	2076
<i>Sergey Redyuk (Technische Universität Berlin)</i>	
Explaining Results of Data-Driven Applications	2081
<i>Nave Frost (Tel Aviv University)</i>	
Towards Explaining the Effects of Data Preprocessing on Machine Learning	2086
<i>Carlos Vladimiro Gonzalez Zelaya (Newcastle University)</i>	
Don't Fear the REAPER: A Framework for Materializing and Reusing Deep-Learning Models	2091
<i>Melanie B. Sigl (Friedrich-Alexander-Universität Erlangen-Nürnberg)</i>	
Knowledge Representation for Emotion Intelligence	2096
<i>Shuo Wang (Renmin University of China)</i>	
Disambiguation and Result Expansion in Keyword Search Over Relational Databases	2101
<i>Niousha Hormozi (University of Athens)</i>	
Event Recommendation using Social Media	2106
<i>Sreekanth Madisetty (Indian Institute of Technology Hyderabad)</i>	

TKDE Posters

Location Inference for Non-Geotagged Tweets in User Timelines [Extended Abstract]	2111
<i>Pengfei Li (Zhejiang University, China), Hua Lu (Aalborg University), Nattiya Kanhabua (NTENT, Barcelona, Spain), Sha Zhao (Zhejiang University), and Gang Pan (Zhejiang University)</i>	
Efficient Parallel Skyline Query Processing for High-Dimensional Data	2113
<i>Mingjie Tang (Ant Financial), Yongyang Yu (Purdue University), Walid G. Aref (Purdue University), Qutaibah M. Malluhi (Qatar University), and Mourad Ouzzani (Qatar Computing Research Institute)</i>	
On Generalizing Collective Spatial Keyword Queries (Extended Abstract)	2115
<i>Harry Kai-Ho Chan (Hong Kong University of Science and Technology), Cheng Long (Nanyang Technological University), and Raymond Chi-Wing Wong (Hong Kong University of Science and Technology)</i>	
A Novel Representation and Compression for Queries on Trajectories in Road Networks (Extended Abstract)	2117
<i>Xiaochun Yang (Northeastern University, China), Bin Wang (Northeastern University, China), Kai Yang (Northeastern University, China), Chengfei Liu (Swinburne University of Technology), and Baihua Zheng (Singapore Management University)</i>	

Efficient Multi-Class Probabilistic SVMs on GPUs	2119
<i>Zeyi Wen (National University of Singapore), Jiashuai Shi (South China University of Technology), Bingsheng He (National University of Singapore), Jian Chen (South China University of Technology), and Yawen Chen (South China University of Technology)</i>	
C2Net: A Network-Efficient Approach to Collision Counting LSH Similarity Join(Extended Abstract)	2121
<i>Hangyu Li (City University of Hong Kong), Sarana Nutanong (Vidyasirimedhi Institute of Science and Technology), Hong Xu (City University of Hong Kong), Chenyun Yu (City University of Hong Kong), and Foryu Ha (City University of Hong Kong)</i>	
LinkBlackHole*: Robust Overlapping Community Detection Using Link Embedding (Extended Abstract)	2123
<i>Jungeun Kim (KAIST), Sungsu Lim (Chungnam National University), Jae-Gil Lee (KAIST), and Byung Suk Lee (University of Vermont)</i>	
Fusion OLAP: Fusing the Pros of MOLAP and ROLAP Together for In-memory OLAP (Extended Abstract)	2125
<i>Yansong Zhang (Renmin University of China), Yu Zhang (National Satellite Meteorological Centre, Beijing, China), Shan Wang (Renmin University of China), and Jiaheng Lu (University of Helsinki)</i>	
In Search of Indoor Dense Regions: An Approach Using Indoor Positioning Data	2127
<i>Huan Li (Zhejiang University), Hua Lu (Aalborg University), Lidan Shou (Zhejiang University), Gang Chen (Zhejiang University), and Ke Chen (Zhejiang University)</i>	
Optimizing Quality for Probabilistic Skyline Computation and Probabilistic Similarity Search (Extended Abstract)	2129
<i>Xiaoye Miao (Zhejiang University), Yunjun Gao (Zhejiang University), Linlin Zhou (Zhejiang University), Wei Wang (University of New South Wales), and Qing Li (The Hong Kong Polytechnic University)</i>	
On Efficiently Answering Why-Not Range-Based Skyline Queries in Road Networks (Extended Abstract)	2131
<i>Xiaoye Miao (Zhejiang University), Yunjun Gao (Zhejiang University), Su Guo (Zhejiang University), and Gang Chen (Zhejiang University)</i>	
SLADE: A Smart Large-Scale Task Decomposer in Crowdsourcing	2133
<i>Yongxin Tong (Beihang University), Lei Chen (Hong Kong University of Science and Technology), Zimu Zhou (ETH Zurich), Hosagrahar Visvesvaraya Jagadish (University of Michigan), Lidan Shou (Zhejiang University), and Weifeng Lv (Beihang University)</i>	
XINA: Explainable Instance Alignment using Dominance Relationship (Extended Abstract)	2135
<i>Jinyoung Yeo (SK T-Brain), Haeju Park (Yonsei University), Sanghoon Lee (NAVER Corporation), Eric Wonhee Lee (Emory University), and Seung-won Hwang (Yonsei University)</i>	
A Hardware-Accelerated Solution for Hierarchical Index-Based Merge-Join(Extended Abstract)	2137
<i>Zimeng Zhou (City University of Hong Kong), Chenyun Yu (City University of Hong Kong), Sarana Nutanong (Vidyasirimedhi Institute of Science and Technology), Yufei Cui (City University of Hong Kong), Chenchen Fu (City University of Hong Kong), and Chun Jason Xue (City University of Hong Kong)</i>	
Finding Most Popular Indoor Semantic Locations Using Uncertain Mobility Data	2139
<i>Huan Li (Zhejiang University), Hua Lu (Aalborg University), Lidan Shou (Zhejiang University), Gang Chen (Zhejiang University), and Ke Chen (Zhejiang University)</i>	

Uncertain Graph Sparsification (Extended Abstract)	2141
<i>Panos Parchas (Amazon Web Services), Nikolaos Papailiou (NTUA), Dimitris Papadias (HKUST), and Francesco Bonchi (ISI Foundation & Eurecat)</i>	
Rule-Based Entity Resolution on Database with Hidden Temporal Information (Extended Abstract)	2143
<i>Hongzhi Wang (Harbin Institute of Technology), Xiaou Ding (Harbin Institute of Technology), Jianzhong Li (Harbin Institute of Technology), and Hong Gao (Harbin Institute of Technology)</i>	
BRIGHT - Drift-Aware Demand Predictions for Taxi Networks (Extended Abstract)	2145
<i>Amal Saadallah (TU Dortmund), Luis Moreira-Matias (Kreditech Holding SSL, Hamburg, Germany), Ricardo Sousa (University of Porto, Portugal), Jihed Khiari (NEC Laboratories Europe, Germany), Erik Jenelius (KTH Royal Institute of Technology, Sweden), and Joao Gama (University of Porto, Portugal)</i>	
Order-Sensitive Imputation for Clustered Missing Values (Extended Abstract)	2147
<i>Qian Ma (Dalian Maritime University, China), Yu Gu (Northeastern University, China), Wang-Chien Lee (The Pennsylvania State University), and Ge Yu (Northeastern University, China)</i>	
DeepDirect: Learning Directions of Social Ties with Edge-Based Network Embedding (Extended Abstract)...	2149
<i>Chaokun Wang (Tsinghua University), Changping Wang (Tsinghua University), Zheng Wang (Tsinghua University), Xiaojun Ye (Tsinghua University), Jeffrey Xu Yu (Chinese University of Hong Kong), and Bin Wang (Tsinghua University)</i>	
A Utility-Optimized Framework for Personalized Private Histogram Estimation (Extended Abstract)	2151
<i>Yiwen Nie (University of Science and Technology of China), Wei Yang (University of Science and Technology of China), Liusheng Huang (University of Science and Technology of China), Xike Xie (University of Science and Technology of China), Zhenhua Zhao (University of Science and Technology of China), and Shaowei Wang (University of Science and Technology of China)</i>	
Near-Accurate Multiset Reconciliation (Extended Abstract)	2153
<i>Lailong Luo (National University of Defense Technology), Deke Guo (National University of Defence Technology), Xiang Zhao (National University of Defense Technology), Jie Wu (Temple University), Ori Rottenstreich (Orbs Research and Technion), and Xueshan Luo (National University of Defense Technology)</i>	
Answering Why-Not Group Spatial Keyword Queries (Extended Abstract)	2155
<i>Bolong Zheng (Huazhong University of Science and Technology), Kai Zheng (University of Electronic Science and Technology of China), Christian S. Jensen (Aalborg University), Nguyen Quoc Viet Hung (Griffith University), Han Su (University of Electronic Science and Technology of China), Guohui Li (Huazhong University of Science and Technology), and Xiaofang Zhou (The University of Queensland)</i>	
Effective and Efficient Community Search Over Large Directed Graphs (Extended Abstract)	2157
<i>Yixiang Fang (University of Hong Kong), Zhongran Wang (Harbin Institute of Technology), Reynold Cheng (The University of Hong Kong), Hongzhi Wang (Harbin Institute of Technology), and Jiafeng Hu (University of Hong Kong)</i>	

Exploring Communities in Large Profiled Graphs (Extended Abstract)	2159
<i>Yankai Chen (The University of Hong Kong), Yixiang Fang (The University of Hong Kong), Reynold Cheng (The University of Hong Kong), Yun Li (Nanjing University), Xiaojun Chen (Shenzhen University), and Jie Zhang (Nanyang Technological University)</i>	
Index-Based Densest Clique Percolation Community Search in Networks (Extended Abstract)	2161
<i>Long Yuan (Nanjing University of Science and Technology), Lu Qin (QCIS, University of Technology, Sydney), Wenjie Zhang (The University of New South Wales), Lijun Chang (The University of Sydney), and Jianye Yang (Central South University)</i>	
Author Index	2163