



Proceedings of the VLDB Endowment

Volume 16, No. 11 – July 2023

Editors in Chief:

Georgia Koutrika and Jun Yang

Associate Editors:

Alkis Simitsis, Amol Deshpande, Angela Bonifati, Ashwin Machanavajjhala, Badrish Chandramouli,
Boris Glavic, Ce Zhang, Cyrus Shahabi, Dan Olteanu, Eric Lo, Evaggelia Pitoura, Evimaria Terzi,
Gustavo Alonso, Helen (Zi) Huang, Hong Cheng, Kenneth Ross, Khuzaima Daudjee, Kyuseok Shim,
Letizia Tanca, Lucian Popa, Magdalena Balazinska, Meihui Zhang, Neoklis Polyzotis,
Nesime Tatbul, Nikos Mamoulis, Rachel Pottinger, Wenjie Zhang, Wolfgang Gatterbauer,
Wook-Shin Han, Xiaokui Xiao, Yannis Velegrakis, Yanyan Shen, Yi Chen, Yongxin Tong, Zhifeng Bao

Publication Editors:

Manos Athanassoulis, Kostas Stefanidis, Ju Fan

PVLDB – Proceedings of the VLDB Endowment

Volume 16, No. 11, July 2023.

All papers published in this issue will be presented at the 49th International Conference on Very Large Data Bases, Vancouver, Canada, 2023.

Copyright 2023 VLDB Endowment

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org.

Volume 16, Number 11, July 2023

Pages i – ix and 2686 - 3474

ISSN 2150-8097

Available at: <http://www.pvldb.org> and <https://dl.acm.org/journal/pvldb>

TABLE OF CONTENTS

Front Matter

Copyright Notice	i
Table of Contents	ii
PVLDB Organization and Review Board – Vol. 16	vi

Research Papers

A Two-Level Signature Scheme for Stable Set Similarity Joins.....	2686
<i>Daniel Schmitt, Daniel Kocher, Nikolaus Augsten, Willi Mann, Alexander Miller</i>	
Scalable Reasoning on Document Stores via Instance-Aware Query Rewriting.....	2699
<i>Olivier Rodriguez, Federico Ulliana, Marie-Laure Mugnier</i>	
EQUI-VOCAL: Synthesizing Queries for Compositional Video Events from Limited User Interactions.....	2714
<i>Enhao Zhang, Maureen Daum, Dong He, Brandon Haynes, Ranjay Krishna, Magdalena Balazinska</i>	
Lotan: Bridging the Gap between GNNs and Scalable Graph Analytics Engines	2728
<i>Yuhao Zhang, Arun Kumar</i>	
Epoxy: ACID Transactions Across Diverse Data Stores	2742
<i>Peter Kraft, Qian Li, Xinjing Zhou, Peter Bailis, Michael Stonebraker, Xiangyao Yu, Matei Zaharia</i>	
Analyzing Vectorized Hash Tables Across CPU Architectures.....	2755
<i>Maximilian Böther, Lawrence Benson, Ana Klimovic, Tilmann Rabl</i>	
Exploiting Cloud Object Storage for High-Performance Analytics.....	2769
<i>Dominik Durner, Viktor Leis, Thomas Neumann</i>	
A Randomized Blocking Structure for Streaming Record Linkage.....	2783
<i>Dimitrios Karapiperis, Christos Tjortjis, Vassilios S. Verykios</i>	
REmatch: a novel regex engine for finding all matches	2792
<i>Cristian Riveros, Nicolás Van Sint Jan, Domagoj Vrgoč</i>	
ADOPT: Adaptively Optimizing Attribute Orders for Worst-Case Optimal Join Algorithms via Reinforcement Learning.....	2805
<i>Junxiong Wang, Immanuel Trummer, Ahmet Kara, Dan Olteanu</i>	
Triangular Stability Maximization by Influence Spread over Social Networks.....	2818
<i>Zheng Hu, Weiguo Zheng, Xiang Lian</i>	
CORE-Sketch: On Exact Computation of Median Absolute Deviation with Limited Space.....	2832
<i>Haoquan Guan, Ziling Chen, Shaoxu Song</i>	
Fast Search-By-Classification for Large-Scale Databases Using Index-Aware Decision Trees and Random Forests	2845
<i>Christian Lülfi, Denis Mayr Lima Martins, Marcos Antonio Vaz Salles, Yongluan Zhou, Fabian Gieseke</i>	
Semi-Oblivious Chase Termination for Linear Existential Rules: An Experimental Study	2858
<i>Marco Calautti, Mostafa Milani, Andreas Pieris</i>	
Analyzing the Impact of Cardinality Estimation on Execution Plans in Microsoft SQL Server.....	2871
<i>Kukjin Lee, Anshuman Dutt, Vivek Narasayya, Surajit Chaudhuri</i>	

WALTZ: Leveraging Zone Append to Tighten the Tail Latency of LSM Tree on ZNS SSD	2884
<i>Jongsung Lee, Donguk Kim, Jae W. Lee</i>	
Accelerating Aggregation Queries on Unstructured Streams of Data	2897
<i>Matthew Russo, Tatsunori Hashimoto, Daniel Kang, Yi Sun, Matei Zaharia</i>	
QueryBooster: Improving SQL Performance Using Middleware Services for Human-Centered Query Rewriting	2911
<i>Qiushi Bai, Sadeem Alsudais, Chen Li</i>	
Consistent Range Approximation for Fair Predictive Modeling	2925
<i>Jiongli Zhu, Sainyam Galhotra, Nazanin Sabri, Babak Salimi</i>	
SUREL+: Moving from Walks to Sets for Scalable Subgraph-based Graph Representation Learning	2939
<i>Haoteng Yin, Muhan Zhang, Jianguo Wang, Pan Li</i>	
Estimating Single-Node PageRank in $\tilde{O}(\min\{d_t, \sqrt{m}\})$ Time	2949
<i>Hanzhi Wang, Zhewei Wei</i>	
Simple Adaptive Query Processing vs. Learned Query Optimizers: Observations and Analysis.....	2962
<i>Yunjia Zhang, Yannis Chronis, Jignesh M. Patel, Theodoros Rekatsinas</i>	
BP-tree: Overcoming the Point-Range Operation Tradeoff for In-Memory B-trees	2976
<i>Helen Xu, Amanda Li, Brian Wheatman, Manoj Marneni, Prashant Pandey</i>	
HENCE-X: Toward Heterogeneity-agnostic Multi-level Explainability for Deep Graph Networks	2990
<i>Ge Lv, Chen Jason Zhang, Lei Chen</i>	
Automatic Road Extraction with Multi-Source Data Revisited: Completeness, Smoothness and Discrimination	3004
<i>Haitao Yuan, Sai Wang, Zhifeng Bao, Shangguang Wang</i>	
Asymptotically Better Query Optimization Using Indexed Algebra.....	3018
<i>Philipp Fent, Guido Moerkotte, Thomas Neumann</i>	
Normalizing Property Graphs	3031
<i>Philipp Skavantzios, Sebastian Link</i>	
A Deep Dive into Common Open Formats for Analytical DBMSs	3044
<i>Chunwei Liu, Anna Pavlenko, Matteo Interlandi, Brandon Haynes</i>	
Saibot: A Differentially Private Data Search Platform.....	3057
<i>Zezhou Huang, Jiexiang Liu, Daniel Gbenga Alabi, Raul Castro Fernandez, Eugene Wu</i>	
JoinBoost: Grow Trees Over Normalized Data Using Only SQL.....	3071
<i>Zezhou Huang, Rathijit Sen, Jiexiang Liu, Eugene Wu</i>	
R ³ : Record-Replay-Retroaction for Database-Backed Applications.....	3085
<i>Qian Li, Peter Kraft, Michael Cafarella, Çağatay Demiralp, Goetz Graefe, Christos Kozyrakis, Michael Stonebraker, Lalith Suresh, Xiangyao Yu, Matei Zaharia</i>	
Self-Training for Label-Efficient Information Extraction from Semi-Structured Web-Pages	3098
<i>Ritesh Sarkhel, Binxuan Huang, Colin Lockard, Prashant Shiralkar</i>	
Efficient Non-Learning Similar Subtrajectory Search.....	3111
<i>Jiabao Jin, Peng Cheng, Lei Chen, Xuemin Lin, Wenjie Zhang</i>	

Frequency-revealing attacks against Frequency-hiding Order-preserving Encryption	3124
<i>Xinle Cao, Jian Liu, Yongsheng Shen, Xiaohua Ye, Kui Ren</i>	
Efficient Fault Tolerance for Recommendation Model Training via Erasure Coding	3137
<i>Tianyu Zhang, Kaige Liu, Jack Kosaian, Juncheng Yang, Rashmi Vinayak</i>	
SlabCity: Whole-Query Optimization using Program Synthesis	3151
<i>Rui Dong, Jie Liu, Yuxuan Zhu, Cong Yan, Barzan Mozafari, Xinyu Wang</i>	
Scaling Up Structural Clustering to Large Probabilistic Graphs Using Lyapunov Central Limit Theorem.....	3165
<i>Joseph Howie, Venkatesh Srinivasan, Alex Thomo</i>	
Epistemic Parity: Reproducibility as an Evaluation Metric for Differential Privacy	3178
<i>Lucas Rosenblatt, Bernease Herman, Anastasia Holovenko, Wonkwon Lee, Joshua Loftus, Elizabeth Mckinnie, Taras Rumezhak, Andrii Stadnik, Bill Howe, Julia Stoyanovich</i>	
POEM: Pattern-Oriented Explanations of Convolutional Neural Networks	3192
<i>Vargha Dadvar, Lukasz Golab, Divesh Srivastava</i>	
gCore: Exploring Cross-layer Cohesiveness in Multi-layer Graphs.....	3201
<i>Dandan Liu, Zhaonian Zou</i>	
Sieve: A Learned Data-Skipping Index for Data Analytics.....	3214
<i>Yulai Tong, Jiazhen Liu, Hua Wang, Ke Zhou, Rongfeng He, Qin Zhang, Cheng Wang</i>	
Out-of-Order Sliding-Window Aggregation with Efficient Bulk Evictions and Insertions.....	3227
<i>Kanat Tangwongsan, Martin Hirzel, Scott Schneider</i>	
k-Best Egalitarian Stable Marriages for Task Assignment	3240
<i>Siyuan Wu, Leong Hou U, Panagiotis Karras</i>	
Federated Calibration and Evaluation of Binary Classifiers.....	3253
<i>Graham Cormode, Igor L. Markov</i>	
FlashAlloc: Dedicating Flash Blocks By Objects	3266
<i>Jonghyeok Park, Soyee Choi, Gihwan Oh, Soojun Im, Moon-Wook Oh, Sang-Won Lee</i>	
Through the Fairness Lens: Experimental Analysis and Evaluation of Entity Matching.....	3279
<i>Nima Shahbazi, Nikola Danevski, Fatemeh Nargesian, Abolfazl Asudeh, Divesh Srivastava</i>	
Check Out the Big Brain on BRAD: Simplifying Cloud Data Processing with Learned Automated Data Meshes	3293
<i>Tim Kraska, Tianyu Li, Samuel Madden, Markos Markakis, Amadou Ngom, Ziniu Wu, Geoffrey X. Yu</i>	
How Large Language Models Will Disrupt Data Management.....	3302
<i>Raul Castro Fernandez, Aaron J. Elmore, Michael J. Franklin, Sanjay Krishnan, Chenhao Tan</i>	
FASTgres: Making Learned Query Optimizer Hinting Effective	3310
<i>Lucas Woltmann, Jerome Thiessat, Claudio Hartmann, Dirk Habich, Wolfgang Lehner</i>	
Write-Aware Timestamp Tracking: Effective and Efficient Page Replacement for Modern Hardware	3323
<i>Demian Vöhringer, Viktor Leis</i>	

Tigger: A Database Proxy That Bounces With User-Bypass.....	3335
<i>Matthew Butrovich, Karthik Ramanathan, John Rollinson, Wan Shen Lim, William Zhang, Justine Sherry, Andrew Pavlo</i>	
Equitable Data Valuation Meets the Right to Be Forgotten in Model Markets.....	3349
<i>Haocheng Xia, Jinfei Liu, Jian Lou, Zhan Qin, Kui Ren, Yang Cao, Li Xiong</i>	
TSM-Bench: Benchmarking Time Series Database Systems for Monitoring Applications.....	3363
<i>Abdelouahab Khelifati, Mourad Khayati, Anton Dignös, Djellel Difallah, Philippe Cudré-Mauroux</i>	
Cross Modal Data Discovery over Structured and Unstructured Data Lakes	3377
<i>Mohamed Y. Eltabakh, Mayuresh Kunjir, Ahmed K. Elmagarmid, Mohammad Shahmeer Ahmad</i>	
Auto-Tables: Synthesizing Multi-Step Transformations to Relationalize Tables without Using Examples	3391
<i>Peng Li, Yeye He, Cong Yan, Yue Wang, Surajit Chaudhuri</i>	
Effective Entity Augmentation By Querying External Data Sources.....	3404
<i>Christopher Buss, Jasmin Mousavi, Mikhail Tokarev, Arash Termehchy, David Maier, Stefan Lee</i>	
Choose Wisely: An Extensive Evaluation of Model Selection for Anomaly Detection in Time Series.....	3418
<i>Emmanouil Sylligardos, Paul Boniol, John Paparrizos, Panos Trahanias, Themis Palpanas</i>	
Similarity search in the blink of an eye with compressed indices	3433
<i>Cecilia Aguerrebere, Ishwar Singh Bhati, Mark Hildebrand, Mariano Tepper, Theodore Willke</i>	
On Data-Aware Global Explainability of Graph Neural Networks	3447
<i>Ge Lv, Lei Chen</i>	
Declarative Sub-Operators for Universal Data Processing.....	3461
<i>Michael Jungmair, Jana Giceva</i>	

PVLDB ORGANIZATION AND REVIEW BOARD - Vol. 16

Editors in Chief of PVLDB

Georgia Koutrika (Athena Research Center)
Jun Yang (Duke University)

Associate Editors of PVLDB

Alkis Simitsis (Athena Research Center)
Amol Deshpande (University of Maryland at College Park)
Angela Bonifati (Lyon 1 University)
Ashwin Machanavajjhala (Duke University/Tumult Labs)
Badrish Chandramouli (Microsoft Research)
Boris Glavic (Illinois Institute of Technology)
Ce Zhang (ETH Zurich)
Cyrus Shahabi (University of Southern California)
Dan Olteanu (University of Zurich)
Eric Lo (The Chinese University of Hong Kong)
Evaggelia Pitoura (University of Ioannina)
Evmiria Terzi (Boston University)
Gustavo Alonso (ETH Zurich)
Helen (Zi) Huang (University of Queensland)
Hong Cheng (The Chinese University of Hong Kong)
Kenneth Ross (Columbia University)
Khuzaima Daudjee (University of Waterloo)
Kyuseok Shim (Seoul National University)
Letizia Tanca (Politecnico di Milano)
Lucian Popa (IBM Research - Almaden)
Magdalena Balazinska (University of Washington)
Meihui Zhang (Beijing Institute of Technology)
Neoklis Polyzotis (Databricks)
Nesime Tatbul (Intel Labs and MIT)
Nikos Mamoulis (University of Ioannina)
Rachel Pottinger (University of British Columbia)
Wenjie Zhang (University of New South Wales)
Wolfgang Gatterbauer (Northeastern University)

Wook-Shin Han (Pohang University of Science and Technology)
Xiaokui Xiao (National University of Singapore)
Yannis Velegrakis (University of Trento and Utrecht University)
Yanyan Shen (Shanghai Jiao Tong University)
Yi Chen (New Jersey Institute of Technology)
Yongxin Tong (Beihang University)
Zhifeng Bao (RMIT University)

Publication Editors

Manos Athanassoulis (Boston University)
Kostas Stefanidis (Tampere University)
Ju Fan (Renmin University of China)

PVLDB Managing Editor

Wolfgang Lehner (Dresden University of Technology)

PVLDB Advisory Board

Vanessa Braganholo (Universidade Federal Fluminense)
Sourav S Bhowmick (Nanyang Technological University)
Chris Jermaine (Rice University)
Peter Triantafillou (University of Warwick)
Xin Luna Dong (Facebook)
Fatma Ozcan (Google)
Lei Chen (Hong Kong University of S&T)
Juliana Freire (New York University)
Graham Cormode (University of Warwick)
Divesh Srivastava (AT&T Labs-Research)
Wolfgang Lehner (Dresden University of Technology)
Felix Naumann (HPI)
Xuemin Lin (University of New South Wales)
Georgia Koutrika (Athena Research Center)
Jun Yang (Duke University)

Review Board

Abolfazl Asudeh (University of Illinois at Chicago)
Alexander J Ratner (University of Washington)
Alexandra Meliou (University of Massachusetts Amherst)
Amelie Marian (Rutgers University)
Amir Gilad (Duke University)
Amir Shaikhha (University of Edinburgh)
Amrita Roy Chowdhury (University of Wisconsin-Madison)
Anastasios Kementsietsidis (Google Research)
Andrew Crotty (Carnegie Mellon University)
Anna Fariha (Microsoft)
Anton Dignös (Free University of Bozen-Bolzano)
Antonios Deligiannakis (Technical University of Crete)
Arijit Khan (Nanyang Technological University)
Ashraf Aboulnaga (Qatar Computing Research Institute, HBKU)
Asterios Katsifodimos (Delft University of Technology)
Baihua Zheng (Singapore Management University)
Bin Cui (Peking University)
Bingsheng He (National University of Singapore)
Binhang Yuan (ETH Zurich)
Bogdan Cautis (University Paris-Saclay)
Bojan Karlas (ETH Zurich)
Bolin Ding (Data Analytics and Intelligence Lab, Alibaba Group)
Bolong Zheng (Huazhong University of Science and Technology)
Bongki Moon (Seoul National University)
Botong Huang (Alibaba)
Brad Glasbergen (University of Waterloo)
Brandon Haynes (Microsoft Gray Systems Lab)
Cedric Renggli (ETH Zurich)
Chao Zhang (Lyon 1 University)
Chen Li (UC Irvine)
Chengfei Liu (Swinburne University of Technology)
Chengkai Li (The University of Texas at Arlington)
Chengliang Chai (Tsinghua University)
Chong Wang (Amazon)
Cristian Riveros (PUC Chile)
Daichi Amagata (Osaka University)
Dan Kifer (Pennsylvania State University)
Daniel Deutch (Tel Aviv University)
Daniel Kang (Stanford University)
Diego Calvanese (Free University of Bozen)
Dimitrios Skoutas (Athena Research Center)
Dimitris Sacharidis (ULB)
Dirk Habich (TU Dresden)
Dong Deng (Rutgers University - New Brunswick)
Dong Wen (University of New South Wales)
Dong Xie (Pennsylvania State University)
Dongxiang Zhang (Zhejiang University)
Dumitrel Loghin (National University of Singapore)
Elena Ferrari (University of Insubria, Varese)
Eleni Tzirita Zacharatou (IT University of Copenhagen)
Essam Mansour (Concordia University)
Faisal Nawab (University of California at Irvine)
Fan Zhang (Guangzhou University)
Fatemeh Nargesian (University of Rochester)
Fei Chiang (McMaster University)
Floris Geerts (University of Antwerp)
Gao Cong (Nanyang Technological University)

George Fakas (Uppsala University)
George Fletcher (Eindhoven University of Technology the Netherlands)
George Papadakis (University of Athens)
George Papastefanatos (Athena Research Center)
Giovanni Simonini (University of Modena and Reggio Emilia)
Graham Cormode (University of Warwick)
Guna Prasaad (Meta Platforms Inc.)
Guoliang Li (Tsinghua University)
Guoren Wang (Beijing Institute of Technology)
Haibo Hu (Hong Kong Polytechnic University)
Hannes Voigt (Neo4j)
Haridimos Kondylakis (FORTH-ICS)
Holger Pirk (Imperial College)
Huanchen Zhang (Tsinghua University)
Ibrahim Sabek (MIT)
Immanuel Trummer (Cornell University)
Ingo Müller (Google)
James Cheng (The Chinese University of Hong Kong)
Jeffrey Xu Yu (The Chinese University of Hong Kong)
Jens Teubner (TU Dortmund University)
Jia Yu (Washington State University)
Jian Lou (Xidian University)
Jianguo Wang (Purdue University)
Jianliang Xu (Hong Kong Baptist University)
Jianxin Li (Deakin University)
Jiawei Jiang (ETH Zurich)
Jieming Shi (Hong Kong Polytechnic University)
Jinfei Liu (Zhejiang University)
Jing Tang (Hong Kong University of Science and Technology)
John Liagouris (Boston University)
John Paparrizos (University of Chicago)
Joseph Near (University of Vermont)
Junhao Gan (University of Melbourne)
K. Selcuk Candan (Arizona State University)
Kai Wang (University of New South Wales)
Karima Echiabi (Mohammed VI Polytechnic University)
Kartik Nayak (Duke University)
Katja Hose (Aalborg University)
Kexin Rong (Stanford University)
Kun Qian (Amazon)
Kunsoo Park (Seoul National University)
Kyriakos Mouratidis (Singapore Management University)
Laks Lakshmanan (University of British Columbia)
Laurel Orr (Stanford University)
Lei Cao (MIT)
Lei Chen (Hong Kong University of Science and Technology)
Lei Li (Hong Kong University of Science and Technology, Guang Zhou)
Lijun Chang (The University of Sydney)
Lin Ma (Carnegie Mellon University)
Long Yuan (Nanjing University of Science and Technology)
Lu Qin (UTS)
Lucas Lersch (Amazon Web Services)
Lukasz Golab (University of Waterloo)
Matteo Interlandi (Microsoft)
Matteo Lissandrini (Aalborg University)
Matthias Renz (University of Kiel)
Matthias Weidlich (Humboldt University of Berlin)

Michael Abebe (University of Waterloo)
 Michael H Boehlen (University of Zurich)
 Michael Hay (Colgate University/Tumult Labs)
 Michael Mathioudakis (University of Helsinki)
 Michal Friedman (ETH)
 Milos Nikolic (University of Edinburgh)
 Mirek Riedewald (Northeastern University)
 Mohamed Sharaf (United Arab Emirates University)
 Mohammad Sadoghi (University of California, Davis)
 Mostafa Milani (The University of Western Ontario)
 Nick Koudas (University of Toronto)
 Nikolaos Tziavelis (Northeastern University)
 Nikolay Yakovets (Eindhoven University of Technology)
 Ning Wang (Beijing Jiaotong University)
 Oliver A Kennedy (University at Buffalo, SUNY)
 Panagiotis Bouros (Johannes Gutenberg University Mainz)
 Panos Kalnis (King Abdullah University of Science and Technology)
 Panos Vassiliadis (University of Ioannina)
 Paolo Papotti (EURECOM, France)
 Periklis Andritsos (University of Toronto)
 Prashant Pandey (University of Utah)
 Raghav Kaushik (Microsoft)
 Rainer Gemulla (University of Mannheim)
 Raul Castro Fernandez (University of Chicago)
 Raymond Chi-Wing Wong (Hong Kong University of Science and Technology)
 Renata Borovica-Gajic (University of Melbourne)
 Reynold Cheng (The University of Hong Kong, China)
 Riccardo Torlone (Roma Tre University)
 Ronghua Li (Beijing Institute of Technology)
 Ryan C Marcus (MIT)
 Ryan Stutsman (University of Utah)
 Sai Wu (Zhejiang Univ)
 Sairam Gurajada (Apple)
 Sebastian Link (University of Auckland)
 Senjuti Basu Roy (New Jersey Institute of Technology)
 Seokki Lee (University of Cincinnati)
 Shantanu Sharma (New Jersey Institute of Technology)
 Shaoxu Song (Tsinghua University)
 Shiyu Yang (Guangzhou University)
 Shuai Ma (Beihang University)
 Sibo Wang (The Chinese University of Hong Kong)
 Siqiang Luo (Nanyang Technological University)
 Sourav S Bhowmick (Nanyang Technological University)
 Spyros Blanas (The Ohio State University)
 Srikanta Bedathur (IIT Delhi)
 Stefania Dumbrava (ENSIIE)
 Stefano Paraboschi (Universita' degli Studi di Bergamo)
 Steffen Zeuch (DFKI Berlin)
 Steven E Whang (KAIST)
 Stijn Vansummeren (Hasselt University)
 Sudipto Das (Amazon Web Services)
 Tarique Siddiqui (Microsoft Research)
 Theodore Dalamagas (Athena Research Center)
 Thomas Neumann (TU Munich)
 Tian Li (Carnegie Mellon University)
 Tianhao Wang (University of Virginia)
 Tianzheng Wang (Simon Fraser University)
 Tien Tuan Anh Dinh (Singapore University of Technology and Design)
 Torben Bach Pedersen (Aalborg University)
 Utku Sirin (Harvard University)
 Vasiliki Kalavri (Boston University)
 Vassilios S Verykios (Hellenic Open University)
 Walid G Aref (Purdue)
 Wang-Chiew Tan (Facebook AI)
 Weiguo Zheng (Fudan University)
 Wendy Hui Wang (Stevens Institute of Technology)
 Wentao Wu (Microsoft Research)
 Xi He (University of Waterloo)
 Xiang Lian (Kent State University)
 Xiangmin Zhou (RMIT University)
 Xiangyao Yu (University of Wisconsin-Madison)
 Xiaochun Yang (Northeastern University)
 Xiaofei Zhang (University of Memphis)
 Xiaoyang Wang (University of New South Wales)
 Xin Cao (University of New South Wales)
 Xin Huang (Hong Kong Baptist University)
 Yan Zhao (Aalborg University)
 Yang Cao (Kyoto University)
 Yao Lu (Microsoft Research)
 Ye Yuan (Beijing Institute of Technology)
 Yeye He (Microsoft Research)
 Yi Yu (NII)
 Yinghui Wu (Case Western Reserve University)
 Yingxia Shao (BUPT)
 Yixiang Fang (The Chinese University of Hong Kong, Shenzhen)
 Yongluan Zhou (University of Copenhagen)
 You Peng (University of New South Wales)
 You Wu (Google)
 Yufei Tao (The Chinese University of Hong Kong)
 Yuncheng Wu (National University of Singapore)
 Yuyu Luo (Tsinghua University)
 Zeke Wang (Zhejiang University)
 Zhiwei Zhang (Beijing Institute of Technology)
 Zhongle Xie (Zhejiang University)
 Zhuoyue Zhao (University at Buffalo - SUNY)
 Ziawasch Abedjan (Leibniz University Hannover)
 Zimu Zhou (Singapore Management University)

LETTER FROM THE EDITORS IN CHIEF

It is our pleasure to present the eleventh issue of PVLDB (Proceedings of the VLDB) Volume 16. This is the last issue that includes research papers that will be presented at VLDB 2023 in Vancouver.

PVLDB presents original research papers on a broad range of topics related to all aspects of data and information management, spanning from theoretical foundations, system architectures, models and techniques, to novel applications as well as large-scale deployment and evaluation. There are four equally important categories of papers in the research track: (a) regular research, (b) scalable data science, (c) experiment, analysis & benchmark, and (d) vision. Each paper is evaluated by at least three reviewers and an Associate Editor, who summarizes in a meta-review all reviews and the results of a three-week discussion phase during which the reviewers exchange their viewpoints and converge to a joint decision.

This issue, the biggest so far in this volume, includes 60 papers, spanning the topics of: Data Mining and Analytics; Data Privacy and Security; Database Engines; Database Performance and Manageability; Distributed Database Systems; Graph and Network Data; Information Integration and Data Quality; Languages; Machine Learning, AI, and Databases; Novel Database Architectures; Provenance and Workflows; Specialized and Domain-Specific Data Management; Text and Semi-Structured Data; and User Interfaces (these cover all top-level topics of interest defined by our Call for Contributions). The most popular topics in this issue are:

- Database Engines (23 papers);
- Machine Learning, AI, and Databases (19 papers);
- Graph and Network Data (10 papers).

The breakdown according to the paper categories is as follows:

- Scalable data data science (3 papers);
- Experiment, analysis & benchmark (8 papers);
- Vision (2 papers);
- Regular research papers (47 papers).

Out of the 60 papers, 53 were accepted after revision, and 7 were accepted after revision plus shepherding.

This issue is the result of all the work put in by the authors as well as the great commitment and effort of our associate editors and reviewers as well as our proceedings chairs.

Georgia Koutrika and Jun Yang
Editors-in-Chief of PVLDB Vol. 16
Program Chairs for VLDB 2023