



Proceedings of the VLDB Endowment

Volume 19, No. 7 – March 2026

Editors in Chief:

Yanlei Diao and Xiaokui Xiao

Associate Editors:

Alkis Polyzotis, Alkis Simitsis, Andrew Pavlo, Angela Bonifati, Anh Dinh, Antonis Deligiannakis, Ashraf Aboulnaga, Badrish Chandramouli, Bin Cui, Bogdan Cautis, Ce Zhang, Elena Ferrari, Evaggelia Pitoura, Fatemeh Nargesian, Guoliang Li, Hong Cheng, Immanuel Trummer, Jana Giceva, Jennie Rogers, Jian Pei, Jianliang Xu, Karima Echiabi, Katja Hose, Khuzaima Daudjee, Kunsoo Park, Kyuseok Shim, Laks Lakshmanan, Lei Cao, Lei Chen, Li Xiong, Mauro Sozio, Meihui Zhang, Melanie Herschel, Michael Böhlen, Nan Tang, Niv Dayan, Norman May, Raul Castro Fernandez, Raymond Chi-Wing Wong, Reynold Cheng, Ruben Mayer, Semih Salihoglu, Shimin Chen, Sibor Wang, Sourav S Bhowmick, Spyros Blanas, Steven Whang, Sudipto Das, Umar Farooq Minhas, Wei Wang, Wook-Shin Han, Yannis Velegrakis, Yanyan Shen, Yi Chen, Yuncheng Wu, Zhifeng Bao, Zi Huang

Publication Editors:

Andrea Mauri, Giovanni Simonini, Siqiang Luo, Stefan Halfpap, Subhadeep Sarkar

PVLDB – Proceedings of the VLDB Endowment

Volume 19, No. 7, March 2026.

All papers published in this issue will be presented at the 52nd International Conference on Very Large Data Bases, Boston, MA, USA, 2026.

Copyright 2026 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 19, Number 7, March 2026

Pages i – ix and 1427 - 1687

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. 19	iv

Research Papers

CAPS: Cost-Aware ML Pipeline Selection	1427
<i>Antonios Kontaxakis, Dimitris Sacharidis, Alberto Abelló, Sergi Nadal, Alkis Simitsis</i>	
Scalable GPU Acceleration of Scalar Functions in Analytical Databases: Compilation, Benchmarking, and Optimization.....	1441
<i>Kaushik Rajan, Sampath Rajendra, Momin Al-Ghosien, Nicolas Bruno, Carlo Curino, Matteo Interlandi, Yinan Li, Lukas M. Maas, Craig Peeper, Surajit Chaudhuri, Johannes Gehrke</i>	
Secure Multi-Party Sampling over Joins	1455
<i>Qiyao Luo, Quanqing Xu, Chuanhui Yang</i>	
How to Write to SSDs.....	1469
<i>Bohyun Lee, Tobias Ziegler, Viktor Leis</i>	
I/O Optimizations in Graph-Based Disk-Resident Approximate Nearest Neighbor Search: A Design Space Exploration.....	1484
<i>Liang Li, Shufeng Gong, Yanan Yang, Yiduo Wang, Jie Wu</i>	
SIDLE: Tree-structure Aware Indexes for CXL-based Heterogeneous Memory.....	1499
<i>Haoru Zhao, Mingkai Dong, Fangnuo Wu, Haibo Chen</i>	
Scarf: Self-Adaptive Tuning via Multi-Objective Reinforcement Learning for Apache Flink	1516
<i>Liu Liu, Shenghao Gong, Ziquan Fang, Yunjun Gao</i>	
JHQ: Johnson-Lindenstrauss Enhanced Hierarchical Quantization for High-Dimensional Approximate Nearest Neighbor Search.....	1530
<i>Jiabao Han, Mengxuan Zhang, Goce Trajcevski</i>	
Love-at-First-Sight: First Answers Without the Awkward Silence in Big Knowledge Graphs	1544
<i>Giannis Vassiliou, Haridimos Kondylakis</i>	
Understanding Disclosure Risk in Differential Privacy with Applications to Noise Calibration and Auditing	1558
<i>Patricia Guerra-Balboa, Annika Sauer, Héber H. Arcolezi, Thorsten Strufe</i>	
CEMR: An Effective Subgraph Matching Algorithm with Redundant Extension Elimination	1572
<i>Linglin Yang, Xunbin Su, Lei Zou, Xiangyang Gou, Yinnian Lin</i>	
AQD: Online Adaptive Query Dispatcher for HTAP Databases	1586
<i>Yang Wu, Tongliang Li, Xuanhe Zhou, Jianying Wang, Xinjun Yang, Wenchao Zhou, Chunxiao Xing, Yong Zhang</i>	
Near-Duplicate Text Alignment under Weighted Jaccard Similarity	1600
<i>Yuheng Zhang, Miao Qiao, Zhencan Peng, Dong Deng</i>	

FlareDTDG: Harnessing Temporal Recency for Scalable Discrete-Time Dynamic Graph Training.....	1614
<i>Wenjie Huang, Rui Wang, Jing Cao, Tongya Zheng, Xinyu Wang, Mingli Song, Sai Wu, Chun Chen</i>	
OpenSQL: Data-Efficient Text-to-SQL for Open-Source LLMs via Synthesized Intermediate Supervision	1628
<i>Ruilin Hu, Yuyu Luo, Guoliang Li, Shuangqiao Wu, Yun Luo</i>	
Database Views as Explanations for Relational Deep Learning	1643
<i>Agapi Rissaki, Ilias Fountalis, Wolfgang Gatterbauer, Benny Kimelfeld</i>	
Characterizing Parallel Subgraph Matching Performance: A Systematic Study of Interactions, Scalability, and Enumeration.....	1659
<i>Tao Yu, Zhijie Zhang, Weiguo Zheng, Jeffrey Xu Yu, Qiang Zhou, Chuntao Hong</i>	
OBELISK: Efficient Offline Query Planning with Bayesian Optimization-Informed Language Model Reasoning.....	1674
<i>Zhicheng Pan, Wenwen Sun, Yuanjia Zhang, Terence Purcell, Yu Dong, Chengcheng Yang, Rong Zhang, Xuan Zhou, Jianliang Xu</i>	

PVLDB ORGANIZATION AND REVIEW BOARD - Vol. 19

Editors in Chief of PVLDB

Yanlei Diao (Ecole Polytechnique and Amazon Web Services)
Xiaokui Xiao (National University of Singapore)

Associate Editors of PVLDB

Alkis Polyzotis (Databricks)
Alkis Simitsis (Athena Research Center)
Andrew Pavlo (Carnegie Mellon University)
Angela Bonifati (Univ. of Lyon)
Anh Dinh (Deakin University)
Antonis Deligiannakis (Technical University of Crete)
Ashraf Aboulnaga (University of Texas at Arlington)
Badrish Chandramouli (Microsoft Research)
Bin Cui (Peking University)
Bogdan Cautis (University of Paris-Saclay)
Ce Zhang (University of Chicago)
Elena Ferrari (University of Insubria, Varese)
Evaggelia Pitoura (Univ. of Ioannina)
Fatemeh Nargesian (University of Rochester)
Guoliang Li (Tsinghua University)
Hong Cheng (Chinese University of Hong Kong)
Immanuel Trummer (Cornell University)
Jana Giceva (TU Munich)
Jennie Rogers (Northwestern University)
Jian Pei (Duke University)
Jianliang Xu (Hong Kong Baptist University)
Karima Echihabi (Mohammed VI Polytechnic University)
Katja Hose (TU Wien)
Khuzaima Daudjee (University of Waterloo)
Kunsoo Park (Seoul National University)
Kyuseok Shim (Seoul National University)
Laks Lakshmanan (The University of British Columbia)
Lei Cao (University of Arizona)
Lei Chen (Hong Kong University of Science and Technology)
Li Xiong (Emory University)
Mauro Sozio (Luiss University of Rome)
Meihui Zhang (Beijing Institute of Technology)
Melanie Herschel (Nanyang Technological University)
Michael Böhlen (University of Zurich)
Nan Tang (HKUST (GZ))
Niv Dayan (University of Toronto)
Norman May (SAP SE)
Raul Castro Fernandez (The University of Chicago)
Raymond Chi-Wing Wong (Hong Kong University of Science and Technology)
Reynold Cheng (The University of Hong Kong, China)
Ruben Mayer (University of Bayreuth)
Semih Salihoglu (University of Waterloo)
Shimin Chen (Chinese Academy of Sciences)
Sibo Wang (The Chinese University of Hong Kong)
Sourav S Bhowmick (Nanyang Technological University)
Spyros Blanas (The Ohio State University)
Steven Whang (KAIST)

Sudipto Das (Amazon Web Services)
Umar Farooq Minhas (Apple)
Wei Wang (ByteDance)
Wook-Shin Han (POSTECH)
Yannis Velegrakis (Utrecht University)
Yanyan Shen (Shanghai Jiao Tong University)
Yi Chen (NJIT)
Yuncheng Wu (Renmin University of China)
Zhifeng Bao (The University of Queensland)
Zi Huang (University of Queensland)

Publication Editors

Andrea Mauri (Lyon 1 University)
Giovanni Simonini (University of Modena and Reggio Emilia)
Siqiang Luo (Nanyang Technological University)
Stefan Halfpap (TU Berlin)
Subhadeep Sarkar (Brandeis University)

PVLDB Managing Editor

Jun Yang (Duke University)

PVLDB Advisory Board

Alexandra Meliou (University of Massachusetts Amherst)
Divesh Srivastava (AT&T Labs - Research)
Fatma Ozcan (Google)
Felix Naumann (HPI)
Katja Hose (Technische Universität Wien)
Lei Chen (Hong Kong University of Science and Technology)
Matthias Boehm (Technische Universität Berlin)
Meihui Zhang (Beijing Institute of Technology)
Nesime Tatbul (Intel Labs and MIT)
Sourav S. Bhowmick (Nanyang Technological University)
Themis Palpanas (Universite Paris Cite)
Torsten Grust (University of Tuebingen)
Vanessa Braganholo (Universidade Federal Fluminense)
Wolfgang Lehner (TU Dresden)
Xiaokui Xiao (National University of Singapore)
Xin Luna Dong (Meta)
Yanlei Diao (Ecole Polytechnique, Amazon Web Services)

Review Board

Abhishek Roy (Microsoft)
Ahmed Aly (Google)
Ahmed Eldawy (University of California, Riverside)
Ahmed Mahmood (Google)
Alberto Lerner (Computing Flows GmbH)
Alberto Sonnino (MystenLabs & University College London)
Alexander Artikis (University of Piraeus)
Alexander Böhm (SAP SE)
Alexander Krause (TU Dresden)
Amarnadh Sai Eluri (Google)
Amelie Marian (Rutgers University)
Amine Mhedhbi (Polytechnique Montréal)
Amir Shaikhha (University of Edinburgh)
Amr Magdy (University of California Riverside)
Andrea Mauri (Lyon 1 University)
Andreas Züfle (Emory University)
Antonio Boffa (EPFL)
Arash Termehchy (Oregon State University)
Arijit Khan (Bowling Green State University)
Aristides Gionis (KTH Royal Institute of Technology)
Arkaprava Saha (University of Tsukuba)
Arnab Phani (TU Berlin)
Arnd Christian König (Microsoft)
Avinash Kumar (Google)
Baihua Zheng (Singapore Management University)
Berthold Reinwald (IBM Research - Almaden)
Bin Yang (Aalborg University)
Bo Tang (Southern University of Science and Technology)
Bo Zhao (Aalto University)
Bobbi W Yogatama (NVIDIA)
Bojan Karlaš (Harvard University)
Boris Glavic (University of Illinois Chicago)
Brian Kroth (Microsoft)
Brit Youngmann (Technion)
Calisto Zuzarte (IBM)
Carlo Curino (Microsoft)
Chang Ge (University of Minnesota)
Chao Zhang (University of Waterloo)
Chen Zhang (The Hong Kong Polytechnic University)
Cheng Long (Nanyang Technological University)
Chenghao Lyu (Amazon)
Chengkai Li (The University of Texas at Arlington)
Chengliang Chai (Beijing Institute of Technology)
Chenhao Ma (The Chinese University of Hong Kong, Shenzhen)
Christan Grant (University of Florida)
Christos Doulkeridis (University of Piraeus)
Christos Koutras (New York University)
Chrysanthi Kosyfaki (The University of Hong Kong)
Chuan Lei (Oracle)
Chunwei Liu (MIT CSAIL)
Cindy Chen (University of Massachusetts Lowell)
Cong Yue (National University of Singapore)
Daniel Kang (UIUC)
Daniel Ritter (SAP)
Davood Rafiei (University of Alberta)
Demetris Zeinalipour (University of Cyprus)
Dimitrios Skoutas (Athena Research Center)
Dixin Tang (The University of Texas, Austin)
Domagoj Vrgoč (PUC)
Dominik Durner (CedarDB)
Donatella Firmani (Sapienza University)
Donato Tiano (Università degli Studi di Modena e Reggio Emilia)
Dong Deng (Alibaba Cloud Computing)
Dong Xie (Penn State University)
Dumitrel Loghin (OKX)
Eduard Dragut (Temple Univ.)
Egemen Tanin (University of Melbourne)
El Kindi Rezig (University of Utah)
Eleni Tzirita Zacharitou (Hasso Plattner Institute (HPI))
Eliana Pastor (Politecnico di Torino)
Elisa Bertino (Purdue University)
Enzo Veltri (Università della Basilicata)
Ergute Bao (Mohamed bin Zayed University of Artificial Intelligence)
Eser Kandogan (Megagon Labs)
Essam Mansour (Concordia University)
Farnoush Banaei-Kashani (University of Colorado Denver)
Felix M Schuhknecht (Johannes Gutenberg University Mainz)
Florian Hahn (University of Twente)
Florin Rusu (UC Merced)
Francesco Gullo (University of L'Aquila)
Fuad Jamour (Amazon Web Services)
Gao Cong (Nanyang Technological University)
Garima Gaur (Inria Saclay, Ecole Polytechnique)
Gengrui Zhang (Concordia University)
George Christodoulou (TU Delft)
George Papadakis (University of Athens)
George Papastefanatos (ATHENA Research Center)
Georgios Siachamis (Inria)
Georgios R Theodorakis (Nvidia)
Giorgio Vinciguerra (Università di Pisa)
Giuseppina Andresini (University of Bari Aldo Moro)
Goce Trajcevski (Iowa State University)
Grigorios Loukides (King's College London)
Guido Moerkotte (University of Mannheim)
Guillaume Lachaud (Ecole Polytechnique)
Haibo Hu (Hong Kong Polytechnic University)
Haipeng Dai (Nanjing University)
Hannes Voigt (Neo4j)
Hantian Zhang (Google)
Hanzhi Wang (University of Copenhagen)
Haralampos Gavriilidis (BIFOLD & Technische Universität Berlin)
Haridimos Kondylakis (FORTH-ICS & Computer Science Department, University of Crete)
Harry Kai-Ho Chan (The University of Sheffield)
Hazar Harmouch (University of Amsterdam)
Herald Kllapi (Google Inc)
Hongzhi Wang (Harbin Institute of Technology)
Hua Lu (Aalborg University)
Huanchen Zhang (Tsinghua University)
Huiping Chen (University of Birmingham)
Hyungsoo Jung (Seoul National University)
Hyunjoon Kim (Hanyang University)
Ibrahim Sabek (University of Southern California)

Igor Zablotski (Mysten Labs)
 Iliia Petrov (Reutlingen University)
 Ilias Azizi (Mohammed VI Polytechnic University)
 Ilie Sarpe (KTH Royal Institute of Technology)
 Indrakshi Ray (Colorado State University)
 Ioannis Demertzis (UCSC)
 Jaeyoung Do (Seoul National University)
 Jan Hidders (Birkbeck, University of London)
 Jeeta Ann Chacko (Technical University of Munich)
 Jeffrey Xu Yu (The Hong Kong University of Science and Technology (Guangzhou))
 Jelle Hellings (McMaster University)
 Jesús Camacho-Rodríguez (Microsoft)
 Jia Li (Hong Kong University of Science and Technology)
 Jia Zou (Arizona State University)
 Jialin Ding (Princeton University)
 Jianguo Wang (Purdue University)
 Jiannan Wang (Tsinghua University)
 Jianxin Li (Edith Cowan University)
 Jiawei Jiang (Wuhan University)
 Jieming Shi (The Hong Kong Polytechnic University)
 Jignesh Patel (Carnegie Mellon University)
 Jin Wang (Arizona State University)
 Jiwon Seo (Seoul National University)
 Johannes Pietrzyk (TU Dresden)
 Johes Bater (Tufts University)
 John Paparrizos (The Ohio State University, USA & Aristotle University of Thessaloniki, Greece)
 Jonathan Fürst (ZHAW Zurich University of Applied Sciences)
 Jonghyeok Park (Korea University)
 Joseph Near (University of Vermont)
 Ju Fan (Renmin University of China)
 Junghoon Kim (UNIST)
 Junhao Gan (University of Melbourne)
 K. Selçuk Candan (Arizona State University)
 Kaisong Huang (University of Calgary)
 Kai-Uwe Sattler (TU Ilmenau)
 Kanchan Chowdhury (Marquette University)
 Karim Benouaret (Univ Lyon)
 Katsiaryna Mirylenka (Zalando SE)
 Kaustubh Beedkar (Indian Institute of Technology Delhi)
 Keke HUANG (The University of British Columbia)
 Kenneth A Ross (Columbia University)
 KI HYUN TAE (Samsung Research)
 Kostas Stefanidis (Tampere University)
 Kurt Stockinger (Zurich University of Applied Sciences)
 Kwanghyun Park (Yonsei University)
 Kyoungmin Kim (EPFL)
 Larissa Capobianco Shimomura (Hasselt University)
 Le Gruenwald (The University of Oklahoma)
 Lefteris Kokoris Kogias (MystenLabs)
 Lei Zou (Peking University)
 Leong Hou U (University of Macau)
 Leopoldo Bertossi (Carleton University)
 Liang Liang (EPFL)
 Lingyang Chu (McMaster University)
 Loredana Caruccio (University of Salerno)
 Lu Chen (Zhejiang University)
 Luca Gagliardelli (eCampus University)
 Lucas Braun (Oracle Labs)
 Lukasz Golab (University of Waterloo)
 Madhulika Mohanty (Inria & Institut Polytechnique de Paris)
 Magnus Mueller (AWS)
 Mahdi Esmailoghli (HU Berlin)
 Makoto Onizuka (Osaka University)
 Man Lung Yiu (Hong Kong Polytechnic University)
 Manuel Rigger (National University of Singapore)
 Marco Patella (University of Bologna)
 Markos Markakis (Massachusetts Institute of Technology)
 Martin Prammer (Carnegie Mellon University)
 Matteo Ceccarello (University of Padova)
 Matteo Interlandi (Microsoft)
 Matthaïos Olma (MongoDB)
 Matthew J Perron (Microsoft)
 Matthias Boehm (Technische Universität Berlin)
 Matthias Weidlich (Humboldt-Universität zu Berlin)
 Maximilian E Schüle (University of Bamberg)
 Meichun Hsu (Oracle America)
 Miao Qiao (The University of Auckland)
 Miao Yu (Meta)
 Michael J Cahill (University of Sydney)
 Michael Grossniklaus (University of Konstanz)
 Michael Gubanov (Florida State University)
 Michael J Mior (Rochester Institute of Technology)
 Mohamed Eltabakh (Qatar Foundation)
 Mohamed Mokbel (University of Minnesota - Twin Cities)
 Mohammad Sadoghi (University of California, Davis)
 Mohammad Javad Amiri (Stony Brook University)
 Mohammed J Saeed (Apple)
 Mostafa Milani (The University of Western Ontario)
 Nick Koudas (University of Toronto)
 Nikolay Yakovets (TU Eindhoven)
 Oktie Hassanzadeh (IBM Research)
 Oliver A Kennedy (University at Buffalo, SUNY)
 Oscar Romero (Universitat Politècnica de Catalunya)
 Panagiotis Bouras (Johannes Gutenberg University Mainz)
 Panagiotis Karras (University of Copenhagen)
 Paolo Garza (Politecnico di Torino)
 Paolo Merialdo (Università degli Studi Roma Tre)
 Paolo Missier (University of Birmingham)
 Paolo Papotti (EURECOM)
 Patrick Damme (Technische Universität Berlin)
 Patrick Schäfer (Humboldt-Universität zu Berlin)
 Paul Boniol (Inria, Ecole normale supérieure)
 Peizhi Wu (University of Pennsylvania)
 Peng Peng (Hunan University)
 Periklis Chrysogelos (Oracle)
 Peter M. Fischer (University of Augsburg)
 Pierangela Samarati (Università degli Studi di Milano)
 Prashant Pandey (Northeastern University)
 Protiva Rahman (University of Florida)
 Qingqing Ye (Hong Kong Polytechnic University)
 Qitong Wang (Harvard University)
 Qizhen Zhang (University of Toronto)
 Rada Chirkova (NC State University)
 Rainer Gemulla (Universität Mannheim)

Rana Alotaibi (King Abdulaziz City for Science and Technology (KACST))
 Rana Shahout (Harvard)
 Renata Borovica-Gajic (University of Melbourne)
 Renchi Yang (Hong Kong Baptist University)
 Ria Borrromeo (University of Philippines)
 Riccardo Torlone (Roma Tre University)
 Rihan Hai (TU Delft)
 Ritesh Sarkhel (Amazon)
 Roe Shraga (WPI)
 Romila Pradhan (Purdue University)
 Ruiyuan Li (Chongqing University)
 Ryan Marcus (University of Pennsylvania)
 Ryan McKenna (University of Massachusetts Amherst)
 Sai Wu (Zhejiang University)
 Sara Foresti (Universita' degli Studi di Milano)
 Sarah Kleest-Meißner (Hasselt University)
 Sebastian Link (University of Auckland)
 Sebastian Michel (RPTU Kaiserslautern-Landau)
 Sepanta Zeighami (University of California Berkeley)
 Shantanu Sharma (New Jersey Institute of Technology)
 Shaoux Song (Tsinghua University)
 Sharad Mehrotra (U.C. Irvine)
 Shen Liang (Université Paris Cité)
 Shihabur Chowdhury (Apple)
 Shimin Di (The Hong Kong University of Science and Technology)
 Silviu Maniu (Université Grenoble Alpes)
 Siqiang Luo (Nanyang Technological University)
 Søren Keiser Jensen (Aalborg University)
 Soror Sahri (Université Paris Cité)
 Stavros Maroulis (ATHENA Research Center)
 Stefania Dumbrava (ENSIEE, INRIA, IRIF, Télécom-SudParis)
 Stefano Marchesin (Università di Padova)
 Stefano Paraboschi (Universita' degli Studi di Bergamo)
 Steffen Zeuch (TU Berlin)
 Surajit Chaudhuri (Microsoft)
 Suyash Gupta (University of Oregon)
 Sven Groppe (Universität zu Lübeck/Institute of Information Systems (IFIS))
 Tanzima Hashem (Bangladesh University of Engineering and Technology)
 Tarique Ashraf Siddiqui (Microsoft Corporation)
 Tenindra Abeywickrama (RIKEN Center for Computational Science)
 Theodoros Chondrogiannis (Norwegian University of Science and Technology (NTNU))
 Thomas Neumann (TUM)
 Tianyi Li (Aalborg University)
 Tianzheng Wang (Simon Fraser University)
 Tieying Zhang (Bytedance)
 Ting Yu (MBZUAI)
 Tingjian Ge (University of Massachusetts, Lowell)
 Tong Chen (The University of Queensland)
 Torben Bach Pedersen (Aalborg University)
 Tristan Allard (Univ Rennes, CNRS, IRISA)
 Tsz Nam Chan (Shenzhen University)
 Ulf Leser (Humboldt-Universität zu Berlin)
 Utku Sirin (Harvard University)
 Venkata Vamsikrishna Meduri (IBM Research-Almaden)
 Verena Kantere (University of Ottawa)
 Vikram Nathan (Amazon, Inc.)
 Viktor Leis (Technische Universität München)
 Vivek Narasayya (Microsoft)
 Vivek Shah (Samsung)
 Wang-Chien Lee (Pennsylvania State University, USA)
 Wei Dong (Nanyang Technological University)
 WEI LU (Renmin University of China)
 Weiguo Zheng (Fudan University)
 Wenqi FAN (The Hong Kong Polytechnic University)
 Wentao Wu (Microsoft Research)
 Wentao Zhang (Peking University)
 William Schultz (MongoDB)
 Xiang Lian (Kent State University)
 Xiangyao Yu (University of Wisconsin-Madison)
 Xiangyu Ke (Zhejiang University)
 Xiao Qin (Snowflake)
 Xiao Yan (Wuhan University)
 Xiaofei Zhang (University of Memphis)
 Xiaohui Yu (York University)
 Xiaou Ding (Harbin Institute of Technology)
 Xiaoying Wang (Microsoft)
 Xike Xie (University of Science and Technology of China)
 Xin Huang (Hong Kong Baptist University)
 Xuhao Chen (MSU)
 Xun Jian (Northwestern Polytechnical University)
 Xupeng Miao (Purdue University)
 Yael Amsterdamer (Bar-Ilan University)
 Yan Huang (University of North Texas)
 Yang Cao (Institute of Science Tokyo)
 Yang Wang (The Ohio State University)
 YAO LU (NUS)
 Yesdaulet Izenov (Nazarbayev University)
 Yeye He (Microsoft Research)
 Yifan Wang (University of Hawaii)
 Yihao Ang (National University of Singapore)
 Yin Yang (Hamad bin Khalifa University)
 Yinan Li (Microsoft Research)
 Yinghui Wu (Case Western Reserve University)
 Yingtai Xiao (TikTok)
 Yiwen Zhu (Microsoft)
 Yixiang Fang (School of Data Science, The Chinese University of Hong Kong, Shenzhen)
 Yongjoo Park (UIUC)
 Yongluan Zhou (University of Copenhagen)
 Yongxin Tong (Beihang University)
 Yu Yang (City University of Hong Kong)
 Yuan Qiu (Southeast University)
 Yujie Hui (Google)
 Yuke Wang (Rice University)
 Yuyu Luo (HKUST (GZ))
 Zeyu Ding (Binghamton University)
 Zhao Cao (Remin University of China)
 ZHAOJING LUO (Beijing Institute of Technology)
 Zhengjie Miao (Simon Fraser University)
 Zhewei Wei (Renmin University of China)
 ZHIWEI FAN (Meta)
 Zhongle Xie (Zhejiang University)
 Zhuoyue Zhao (University at Buffalo)
 Zoi Kaoudi (IT University of Copenhagen)

Light Load Review Board

Abdulhakim Qahtan (Utrecht University)
Abolfazl Asudeh (University of Illinois Chicago)
Ahmed El-Roby (Carleton University)
Alberto Abelló (Universitat Politècnica de Catalunya)
Alex Conway (Cornell Tech)
Alexander Erben (Technical University of Munich)
Alexandros Labrinidis (University of Pittsburgh)
Anupam Sanghi (TU Darmstadt)
Ariful Azad (Texas A&M University)
Babak Salimi (University of California at San Diego)
Baotong Lu (Microsoft Research)
Bettina Kemme (McGill University)
C. Seshadhri (UCSC)
Curtis Dyreson (Utah State University)
Danfeng Zhang (Duke University)
Daniel Kocher (University of Salzburg)
Davide Mottin (Aarhus University)
Fabian Panse (University of Augsburg)
Felix Naumann (Hasso Plattner Institute, University of Potsdam)
Haitao Yuan (Nanyang Technological University)
Hangdong Zhao (Microsoft)
Hideyuki Kawashima (Keio University)
Hongyan Chang (NUS)
Hongzhi Yin (The University of Queensland)
Ingo Müller (Google)
Jiaheng Lu (University of Helsinki)
Jinho Lee (Seoul National University)
Juan Sequeda (data.world)
Katia Papakonstantinou (Athens University of Economics and Business)
Kexin Rong (Georgia Institute of Technology)
Konstantinos Mamouras (Rice University)
Kyriakos Mouratidis (Singapore Management University)
Lipeng Wan (Georgia State University)
Manolis Terrovitis (Athena Research Center)
Matteo Brucato (Microsoft Research)
Mayuresh Kunjir (Amazon AWS)
Michele Linardi (CYU)
Mo Sha (Alibaba Group)
Murat Demirbas (MongoDB Research)
Myeongjae Jeon (POSTECH)
Nikos Giatrakos (Technical University of Crete)
Odysseas Papapetrou (TU Eindhoven)
Panagiotis Liakos (Athens University of Economics and Business)
Panagiotis Liakos (University of Athens)
Panos Vassiliadis (University of Ioannina)
Paris Carbone (KTH Royal Institute of Technology)
Paul Groth (University of Amsterdam)
Pierre Bourhis (CNRS)
Quang-Trung Ta (National University of Singapore)
Rathijit Sen (Microsoft)
Remy Wang (University of California, Los Angeles)
Renzo Angles (Universidad de Talca)
Seokki Lee (University of Cincinnati)
Sergi Nadal (Universitat Politècnica de Catalunya)
Sivaprasad Sudhir (MIT)
Song Jiang (University of Texas, Arlington)
Sonia Bergamaschi (Università di Modena e Reggio Emilia)
Subarna Chatterjee (Datastax)
Subhadeep Sarkar (Brandeis University)
Sujaya Maiyya (University of Waterloo)
Supawit Chockchowwat (University of Illinois Urbana-Champaign)
Theodore Johnson (AT&T Labs - Research)
Thomas Fahringer (University of Innsbruck)
Thorsten Papenbrock (Philipps University of Marburg)
Vassilis J. Tsotras (UC Riverside)
Venkatesh Emani (Microsoft)
Wenjia He (University of Michigan)
Wolfgang Gatterbauer (Northeastern University)
Xingguang Chen (National University of Singapore)
Yogesh Simmhan (Indian Institute of Science (IISc), Bangalore)
Yuxin Tang (Rice University)
Zhe Jiang (University of Florida)
Ziawasch Abedjan (BIFOLD/TU Berlin)

LETTER FROM THE EDITORS IN CHIEF

We are pleased to present the seventh issue of Volume 19 of the Proceedings of the VLDB Endowment (PVLDB).

This seventh issue of Volume 19 includes 18 papers spanning the following primary subject areas: Data Privacy and Security (2); Database Engines (3); Graph and Network Data (2); Information Integration and Data Quality (2); Machine Learning, AI, and Databases (5); Novel Database Architectures (2); Specialized and Domain-Specific Data Management (1); and User Interfaces (1). Together, these papers represent a broad and diverse cross-section of the areas addressed by PVLDB.

This issue highlights several research directions that continue to shape the evolution of modern data management. A particularly prominent theme is the increasingly close interaction between database systems, machine learning, and AI-driven data interaction. Several papers explore this connection from complementary perspectives, including cost-aware pipeline selection for automated machine learning, adaptive query dispatching for HTAP systems, data-efficient Text-to-SQL with open-source LLMs, language-model-assisted query planning, and database views as explanations for relational deep learning. Together, these works reflect the expanding role of data systems not only as infrastructure for AI workloads, but also as environments in which AI techniques are becoming integral to optimization, interaction, automation, and explainability.

At the same time, this issue underscores the continued importance of foundational systems research. Several papers revisit core architectural and performance challenges, including GPU acceleration of scalar functions in analytical databases, out-of-place write designs for SSD-based database systems, and tree-structure-aware indexing for CXL-based heterogeneous memory. These contributions illustrate how advances in storage, memory, and hardware acceleration continue to open new opportunities for rethinking long-standing database design principles.

Vector search, high-dimensional data management, and graph analytics also feature prominently in this issue. The papers in this area include work on I/O-efficient disk-resident approximate nearest neighbor search, quantization techniques for scalable high-dimensional ANN, scalable dynamic graph training, efficient subgraph matching, and a systematic study of parallel subgraph matching performance. These contributions reflect the growing importance of specialized data structures and execution strategies for graph, vector, and AI-oriented workloads.

Another notable theme is trustworthy and privacy-preserving data management. This issue includes contributions on secure multi-party sampling over joins and on understanding disclosure risk in differential privacy, with implications for noise calibration and auditing. These papers emphasize that as data systems support increasingly collaborative and sensitive analytics, privacy guarantees, security, and practical efficiency must be advanced together.

The issue also includes work on knowledge graph exploration and near-duplicate text alignment, further illustrating the breadth of information integration, data quality, and text-centric data management challenges addressed by the community. Complementing these technical advances, several papers place strong emphasis on systematic evaluation, benchmarking, design-space exploration, auditing, and performance characterization, underscoring the community's continued commitment to rigor, reproducibility, and trustworthy empirical methodology.

Overall, this issue presents a rich combination of work on AI-enabled data management, hardware-conscious system design, vector and graph processing, privacy-preserving analytics, and rigorous evaluation. We hope that readers will find in these papers both practical advances and new perspectives on the evolving foundations of data management.

We would like to thank the authors for their high-quality submissions, the reviewers and associate editors for their careful and thorough evaluations, and the PVLDB publication team for their continued support in producing this issue.

Yanlei Diao and Xiaokui Xiao
Editors-in-Chief of PVLDB Vol. 19
Program Chairs for VLDB 2026