



# Proceedings of the VLDB Endowment

Volume 17, No. 1 – September 2023

**Editors in Chief:**

Meihui Zhang and Cyrus Shahabi

**Associate Editors:**

Alkis Polyzotis, Amol Deshpande, Angela Bonifati, Ashraf Aboulnaga, Ashwin Machanavajjhala,  
Beng Chin Ooi, Boris Glavic, Ce Zhang, Divy Agrawal, Eric Lo, Fatma Ozcan, Guoliang Li,  
Jeffrey Xu Yu, Jian Pei, Jianliang Xu, Johannes Gehrke, K. Selçuk Candan, Kyuseok Shim, Li Xiong,  
Magdalena Balazinska, Matthias Boehm, Melanie Herschel, Michael Böhlen,  
Nikos Mamoulis, Pinar Tozun, Rachel Pottinger, Sharad Mehrotra, Surajit Chaudhuri, Tamer Özsu,  
Tien Tuan Anh Dinh, Walid Aref, Wei Wang, Xiaokui Xiao, Yanyan Shen, Yongxin Tong, Zi Huang

**Publication Editors:**

Ju Fan, Yang Cao, Xiaou Ding

PVLDB – Proceedings of the VLDB Endowment

Volume 17, No. 1, September 2023.

All papers published in this issue will be presented at the 50th International Conference on Very Large Data Bases, Guangzhou, China, 2024.

## 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](mailto:info@vldb.org).

Volume 17, Number 1, September 2023

Pages i – vii and 1 - 78

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. 17 .....	iii

### Research Papers

DecLog: Decentralized Logging in Non-Volatile Memory for Time Series Database Systems .....	1
<i>Bolong Zheng, Yongyong Gao, Jingyi Wan, Lingsen Yan, Long Hu, Bo Liu, Yunjun Gao, Xiaofang Zhou, Christian S. Jensen</i>	
Efficient Dynamic Weighted Set Sampling and Its Extension .....	15
<i>Fangyuan Zhang, Mengxu Jiang, Sibow Wang</i>	
ZIP: Lazy Imputation during Query Processing .....	28
<i>Yiming Lin, Sharad Mehrotra</i>	
FedGTA: Topology-aware Averaging for Federated Graph Learning.....	41
<i>Xunkai Li, Zhengyu Wu, Wentao Zhang, Yinlin Zhu, Ronghua Li, Guoren Wang</i>	
Host Profit Maximization: Leveraging Performance Incentives and User Flexibility .....	51
<i>Xueqin Chang, Xiangyu Ke, Lu Chen, Congcong Ge, Ziheng Wei, Yunjun Gao</i>	
DP-PQD: Privately Detecting Per-Query Gaps In Synthetic Data Generated By Black-Box Mechanisms .....	65
<i>Shweta Patwa, Danyu Sun, Amir Gilad, Ashwin Machanavajhala, Sudeepa Roy</i>	

## PVLDB ORGANIZATION AND REVIEW BOARD - Vol. 17

### Editors in Chief of PVLDB

Meihui Zhang (Beijing Institute of Technology)  
Cyrus Shahabi (University of Southern California)

### Associate Editors of PVLDB

Alkis Polyzotis (Databricks)  
Amol Deshpande (University of Maryland at College Park)  
Angela Bonifati (Lyon 1 University)  
Ashraf Aboulnaga (Qatar Computing Research Institute, HBKU)  
Ashwin Machanavajjhala (Duke)  
Beng Chin Ooi (NUS)  
Boris Glavic (Illinois Institute of Technology)  
Ce Zhang (ETH)  
Divy Agrawal (University of California, Santa Barbara)  
Eric Lo (Chinese University of Hong Kong)  
Fatma Ozcan (Google)  
Guoliang Li (Tsinghua University)  
Jeffrey Xu Yu (Chinese University of Hong Kong)  
Jian Pei (Simon Fraser University)  
Jianliang Xu (Hong Kong Baptist University)  
Johannes Gehrke (Microsoft)  
K. Selçuk Candan (Arizona State University)  
Kyuseok Shim (Seoul National University)  
Li Xiong (Emory University)  
Magdalena Balazinska (UW)  
Matthias Boehm (Technische Universität Berlin)  
Melanie Herschel (Universität Stuttgart)  
Michael Böhlen (University of Zurich)  
Nikos Mamoulis (University of Ioannina)  
Pinar Tozun (IT University of Copenhagen)  
Rachel Pottinger (Univ. of British Columbia)  
Sharad Mehrotra (U.C. Irvine)  
Surajit Chaudhuri (Microsoft)

Tamer Özsu (University of Waterloo)  
Tien Tuan Anh Dinh (Deakin University)  
Walid Aref (Purdue University)  
Wei Wang (ByteDance)  
Xiaokui Xiao (National University of Singapore)  
Yanyan Shen (Shanghai Jiao Tong University)  
Yongxin Tong (Beihang University)  
Zi Huang (University of Queensland)

### Publication Editors

Ju Fan (Renmin University of China)  
Yang Cao (Hokkaido University)  
Xiaou Ding (Harbin Institute of Technology)

### PVLDB Managing Editor

Wolfgang Lehner (Dresden University of Technology)

### PVLDB Advisory Board

Vanessa Braganholo (Universidade Federal Fluminense)  
Sourav S Bhowmick (Nanyang Technological University)  
Torsten Grust (University of Tuebingen)  
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)  
Felix Naumann (HPI)  
Georgia Koutrika (Athena Research Center)  
Jun Yang (Duke University)  
Meihui Zhang (Beijing Institute of Technology)  
Cyrus Shahabi (University of Southern California)  
Nesime Tatbul (Intel Labs and MIT)  
Themis Palpanas (Universite Paris Cite)

## Review Board

Abolfazl Asudeh (University of Illinois Chicago)  
Aditya Parameswaran (University of California, Berkeley)  
Ahmed S. Abdelhamid (Purdue University)  
Ahmed Eldawy (University of California, Riverside)  
Ahmed El-Roby (Carleton University)  
Ahmed Mahmood (Google)  
Alberto Lerner (University of Fribourg, Switzerland)  
Alexander Thomson (Google)  
Amr Magdy (University of California Riverside)  
Andreas Züfle (Emory University)  
Angelos Christos Anadiotis (Oracle)  
Anja Gruenheid (Microsoft)  
Anthony Tung (National U. of Singapore)  
Anton Dignös (Free University of Bozen-Bolzano, Italy)  
Arijit Khan (Aalborg University)  
Avrilia Floratou (Microsoft)  
Baihua Zheng (Singapore Management University)  
Bailu Ding (Microsoft Research)  
Berthold Reinwald (IBM Research-Almaden)  
Bin Yang (East China Normal University)  
Bingsheng He (National University of Singapore)  
Bolin Ding (Data Analytics and Intelligence Lab, Alibaba Group)  
Brandon Haynes (Microsoft Gray Systems Lab)  
Chao Zhang (University of Waterloo)  
Cheng Long (Nanyang Technological University)  
Chengfei Liu (Swinburne University of Technology)  
Chengkai Li (The University of Texas at Arlington)  
Chengliang Chai (Beijing Institute of Technology)  
Chrysanthi Kosyfaki (University of Ioannina)  
Chunwei Liu (MIT)  
Cong Yan (Microsoft research)  
Daisy Zhe Wang (University of Florida)  
Dan Kifer (Pennsylvania State Univ., USA)  
Dan Lin (Vanderbilt University)  
Daniel Kang (UIUC)  
Demetrios Zeinalipour-Yazti (University of Cyprus)  
Dimitris Papadias (HKUST)  
Dong Deng (Rutgers University - New Brunswick)  
Dong Wen (University of New South Wales)  
Dong Xie (Penn State University)  
Dongxiang Zhang (Zhejiang University)  
Dumitrel Loghin (National University of Singapore)  
Egemen Tanin (University of Melbourne)  
El Kindi Rezig (Massachusetts Institute of Technology)  
Elena Ferrari (University of Insubria, Varese)  
Eser Kandogan (Megagon Labs)  
Essam Mansour (Concordia University)  
Fan Zhang (Guangzhou University)  
Fatemeh Nargesian (University of Rochester)  
Fei Chiang (McMaster University)  
Feng Zhang (Renmin University of China)  
Florin Rusu (UC Merced)  
Gabriel Ghinita (Hamad Bin Khalifa University)  
Gao Cong (Nanyang Technological University)  
George Fakas (Uppsala University)  
Haibo Hu (Hong Kong Polytechnic University)  
Holger Pirk (Imperial College)  
Hong Cheng (Chinese University of Hong Kong)  
Hongzhi Wang (Harbin Institute of Technology)  
Hua Lu (Roskilde University)  
Huanchen Zhang (Tsinghua University)  
Huiping Cao (New Mexico State University)  
Ibrahim Sabek (MIT)  
Ilaria Bartolini (University of Bologna)  
Jana Giceva (TU Munich)  
Jennie Rogers (Northwestern University)  
Jia Zou (Arizona State University)  
Jian Lou (Zhejiang University)  
Jiangshan Yu (Monash University)  
Jianguo Wang (Purdue University)  
Jiannan Wang (Simon Fraser University)  
Jianqiu Xu (Nanjing University of Aeronautics and Astronautics)  
Jianxin Li (Deakin University)  
Jieming Shi (The Hong Kong Polytechnic University)  
Jin Wang (Megagon Labs)  
Jinfei Liu (Zhejiang University)  
Johes Bater (Tufts University)  
John Liagouris (Boston University)  
Jonathan Goldstein (Microsoft)  
Ju Fan (Renmin University of China)  
Juhua Hu (University of Washington)  
Kai Wang (Shanghai Jiao Tong University)  
Kangfei Zhao (Beijing Institute of Technology)  
Karima Echihabi (Mohammed VI Polytechnic University)  
Katja Hose (TU Wien)  
Khuzaima Daudjee (University of Waterloo)  
Kyoungmin Kim (POSTECH)  
Lawrence Benson (HPI, University of Potsdam)  
Lei Chen (Hong Kong University of Science and Technology)  
Lei Zou (Peking University)  
Leong Hou U (University of Macau)  
Lin Ma (University of Michigan)  
Linyang Chu (McMaster University)  
Liyue Fan (UNC Charlotte)  
Lu Chen (Zhejiang University)  
Luigi Bellomarini (Banca d'Italia)  
Madelon Hulsebos (University of Amsterdam)  
Manolis Terrovitis (IMIS, Athena RC)  
Marco Patella (University of Bologna)  
Mario Nascimento (Northeastern University)  
Matteo Lissandrini (Aalborg University)  
Matthias Renz (University of Kiel)  
Michael Hay (Colgate University & Tumult Labs)  
mingjie tang (Ant Financial)  
Mirek Riedewald (Northeastern University)  
Mohamed S. Hassan (Google)  
Mohamed Mokbel (University of Minnesota - Twin Cities)  
Mohammad Javad Amiri (University of Pennsylvania)  
Mostafa Milani (The University of Western Ontario)  
Mourad OUZZANI (Qatar Computing Research Institute, HBKU)  
Nesime Tatbul (Intel Labs and MIT)  
Norman May (SAP SE)  
Oliver A Kennedy (University at Buffalo, SUNY)

Panagiotis Bouros (Johannes Gutenberg University Mainz)  
 Papotti Paolo (EURECOM)  
 Patrick Damme (Technische Universität Berlin)  
 Peng Peng (Hunan University)  
 Philippe Bonnet (IT Univ Copenhagen, Denmark)  
 Pinar Karagoz (METU, Turkey)  
 Prashant Pandey (University of Utah)  
 Primal Pappachan (Penn State University)  
 Qichen Wang (Hong Kong Baptist University)  
 Qing Liu (Zhejiang University)  
 Qun Chen (Northwestern Polytechnical University)  
 Renata Borovica-Gajic (University of Melbourne)  
 Rihan Hai (TU Delft)  
 Ritesh Ahuja (Oracle Labs)  
 Roger Zimmermann (NUS)  
 Ronghua Li (Beijing Institute of Technology)  
 Sai Wu (Zhejiang Univ)  
 Sanjay Krishnan (UChicago)  
 Senjuti Basu Roy (NJIT)  
 Seokki Lee (University of Cincinnati)  
 Shantanu Sharma (New Jersey Institute of Technology)  
 Shaofeng Cai (National University of Singapore)  
 Shaoxu Song (Tsinghua University)  
 Shuai Ma (Beihang University)  
 Shuang Hao (Beijing Jiaotong University)  
 Sibowang (The Chinese University of Hong Kong)  
 Stefania Dumbrava (ENSIIE)  
 Stefano Paraboschi (Universita' degli Studi di Bergamo)  
 Sujaya Maiyya (University of Waterloo)  
 Tarique Siddiqui (Microsoft Research)  
 Thanana Ghanem (Metro State University)  
 Thang Dinh (VCU)  
 Themis Palpanas (Universite Paris Cite)  
 Thomas Neumann (TUM)  
 Tianhao Wang (University of Virginia)  
 Tianzheng Wang (Simon Fraser University)  
 Tieying Zhang (Bytedance)  
 Tristan Allard (Univ Rennes, CNRS, IRISA)  
 Umar Farooq Minhas (Apple)  
 Utku Sirin (Harvard University)  
 Viktor Leis (Technische Universität München)  
 Vincenzo Gulisano (Chalmers University of Technology)  
 Vraj Shah (IBM Research)  
 Wang-Chien Lee (Pennsylvania State University, USA)  
 WEI LU (Renmin University of China)

Wei Wang (Hong Kong University of Science and Technology (Guangzhou))  
 Wei-Shinn Ku (Auburn University)  
 Wenchao Zhou (Alibaba Group)  
 Wendy Hui Wang (Stevens Institute of Technology)  
 Xiang Lian (Kent State University)  
 Xiang Zhao (National University of Defence Technology)  
 Xiangyao Yu (University of Wisconsin-Madison)  
 Xiao Hu (Duke University)  
 Xiao Hu (University of Waterloo)  
 Xiaochun Yang (Northeastern University)  
 Xiaofang Zhou (The Hong Kong University of Science and Technology)  
 Xiaofei Zhang (University of Memphis)  
 Xiaohui Yu (York University)  
 Xiaoli Wang (Xiamen University)  
 Xin Huang (Hong Kong Baptist University)  
 Xin Wang (Tianjin University)  
 Xingquan Zhu (Florida Atlantic University)  
 Yanfeng Zhang (Northeastern University)  
 Yang Cao (Hokkaido University)  
 Yannis Chronis (Google)  
 Yao Lu (Microsoft Research)  
 Ye Yuan (Beijing Institute of Technology)  
 Yeye He (Microsoft Research)  
 Ying Zhang (University of Technology Sydney)  
 Yingxia Shao (BUPT)  
 Yu Yang (City University of Hong Kong)  
 Yuhao Zhang (University of California, San Diego)  
 Yuncheng Wu (National University of Singapore)  
 Yunjun Gao (Zhejiang University)  
 Yuval Moskovitch (Ben Gurion University)  
 Yuxiang Zeng (Beihang University)  
 Zhaojing Luo (National University of Singapore)  
 Zhengjie Miao (Duke University)  
 Zhichao Cao (Arizona State University)  
 Zhifeng Bao (RMIT University)  
 Zhiwei Zhang (Beijing Institute of Technology)  
 Zhongle Xie (Zhejiang University)  
 Zhuoyue Zhao (University at Buffalo - SUNY)  
 Ziawasch Abedjan (Leibniz Universität Hannover)  
 Ziliang Lai (Chinese University of Hong Kong)  
 Zimu Zhou (City University of Hong Kong)

## LETTER FROM THE EDITORS IN CHIEF

We are delighted to introduce you to the inaugural edition of Volume 17 of PVLDB (Proceedings of the VLDB). PVLDB is dedicated to showcasing original research papers that encompass a wide spectrum of subjects within the realm of data and information management. Our coverage spans from fundamental theoretical principles and cutting-edge system architectures to innovative models, techniques, novel applications, and the comprehensive assessment and deployment of large-scale solutions. In our research track, we feature four equally significant categories of papers: (a) regular research, (b) scalable data science (SDS), (c) experiment, analysis & benchmark (EA&B), and (d) vision papers.

PVLDB is committed to providing valuable and constructive feedback through a rigorous review process. All submissions undergo meticulous peer review by a team of accomplished Associate Editors and dedicated reviewers. Each paper receives comprehensive evaluation from a minimum of three reviewers, along with the oversight of an Associate Editor. During a three-week discussion phase, reviewers engage in a thorough exchange of perspectives, ultimately converging on a consensus, which is summarized in a meta-review. Some submissions may proceed to a revision phase, affording authors a three-month window to refine their work for subsequent review cycles. Accepted papers are subsequently published in the journal and proudly presented at the forthcoming VLDB conference.

We have retained most of the processes from the previous year, maintaining continuity and consistency. Similar to last year, a subset of submissions may now proceed through an additional formal shepherding phase, where an assigned shepherd collaborates with the authors to ensure the paper's final version meets the requisite standards. Alongside the continued use of Microsoft's Conference Management Toolkit (CMT) to oversee submissions, reviews, and revisions, we leverage the capabilities of the Toronto Paper Matching System (TPMS) for suggesting review assignments, the Conflict of Interest Detection & Management System (CLOSET) for identifying potential conflicts of interest, and iThenticate for detecting instances of plagiarism.

This year, we have introduced significant enhancements aimed at streamlining the conflicts of interest (COIs) process. Our primary emphasis is on reducing the administrative burden for authors rather than relying solely on COI checks for identification and punitive actions against missing disclosures. We firmly believe that these omissions are often the result of inadvertent oversights. Nevertheless, we maintain robust procedures to identify and address intentional omissions.

Towards this end, we've transitioned from submission-level COIs to individual COIs. Under this new system, each author needs to input their conflicts into CMT just once, regardless of how many papers they submit throughout the month or across multiple months to PVLDB Vol. 17. They only need to update their COIs monthly to account for any potential new conflicts that may have arisen.

Recognizing the challenge of authors ensuring the completeness of their COIs, especially with our extensive review board comprising more than 240 members, we've implemented two key measures. Firstly, immediately after the submission deadline each month, we use our automated tool, CLOSET, to identify potential COIs by scanning DBLP and recent submissions to other database conferences to detect recent co-authorships. We populate these conflicts to CMT for each and every author. Secondly, we provide a 48-hour grace period for authors to review the populated list and add any additional COIs, such as friendships, collaborations without associated papers/submissions, or advising relationships, to ensure comprehensive disclosure. These changes aim to facilitate a more efficient and author-friendly COI process.

However, we kindly request that authors refrain from deleting potential COIs that do not constitute genuine conflicts of interest. Instead, authors can report these instances using an online form provided for this purpose. We routinely review and investigate these deletions, and instances where valid COIs were intentionally removed may lead to the rejection of the paper and other penalties, as determined by the PVLDB advisory board.

Furthermore, the online form also enables authors to report any pre-populated COIs that are not valid conflicts, often referred to as False Positives (FPs). We share these FPs with our colleagues responsible for managing the CLOSET tool to prevent the recurrence of such false positives in authors' future conflict reports. Specifically, if an FP arises due to a name that is identical to another author's, we request that authors include their DBLP URL in the online form to mitigate such occurrences in future assessments. Additionally, if an FP is associated with a community report, we ask authors to provide the title and year of the report in DBLP within the online form. Please note that we verify these inputs on our end. This process not only helps eliminate FPs for future PVLDB submissions but also assists authors in avoiding these FPs in their submissions to other conferences that utilize the CLOSET tool.

Consistent with the previous year, authors are expected to provide supplementary materials, such as code, data, and other implementation components, which were employed in generating the reported results in their paper. In cases where compelling reasons exist, hindering compliance with this availability requirement, authors should furnish an explanation and may be granted an exemption. As part of the meta-review process for accepted submissions, Associate Editors use a

standard rubric to assess the availability of supplemental materials, ensuring their openness and permanence, as well as the readability of instructions for the reuse of the artifacts by other members of the community. Accepted papers that furnish supplementary materials in accordance with the availability requirement receive an official ACM badge.

This first issue of PVLDB's Volume 17 includes six papers, spanning the topics of spatial and temporal databases, probabilistic data, graph and network data management, query processing and data privacy. Out of the six papers, one was a straight accept, and five were accepted after revision. One paper is in the scalable data science (SDS) category, and the rest are regular research papers.

We are very grateful to our board of associate editors and reviewers as well as our proceedings chairs who contribute to the success of PVLDB.

Meihui Zhang and Cyrus Shahabi  
Editors-in-Chief of PVLDB Vol. 17  
Program Chairs for VLDB 2024