Proceedings of the VLDB Endowment

Volume 18, No. 12 – August 2025

PVLDB – Proceedings of the VLDB Endowment

Volume 18, No. 12, August 2025.

All papers published in this issue will be presented at the 51st International Conference on Very Large Data Bases, London, United Kingdom, 2025.

Copyright 2025 VLDB Endowment

Volume 18, Number 12, August 2025

Pages i – xix and 4763 - 5539

ISSN 2150-8097

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

TABLE OF CONTENTS

Front Matter

Industrial Papers

Demonstrations

Tutorials

Panels

PVLDB ORGANIZATION AND REVIEW BOARD - Vol. 18

Editors in Chief of PVLDB
Themis Palpanas (University Paris Cite)
Nesime Tatbul (Intel Labs and MIT)

Associate Editors of PVLDB
Walid G. Aref (Purdue University)
Manos Athanassoulis (Boston University)
Carsten Binnig (Technical University of Darmstadt)
Spyros Blanas (Ohio State University)
Matthias Boehm (Technical University of Berlin)
Angela Bonifati (University of Lille)
K. Selcuk Candan (Arizona State University)
Lei Cao (University of Arizona)
Raul Castro Fernandez (University of Chicago)
Lei Chen (Hong Kong University of Science and Technology)
Shimin Chen (Chinese Academy of Sciences)
Yi Chen (New Jersey Institute of Technology)
Reynold Cheng (University of Hong Kong)
Alvin Cheung, University of California (Berkeley)
Sudipto Das (Amazon Web Services)
Niv Dayan (University of Toronto)
Antonis Deligiannakis (Technical University of Crete)
Jens Dittrich (Saarland University)
Xin Luna Dong (Meta)
Karima Echihabi (Mohammed VI Polytechnic University)
Alan Fekete (University of Sydney)
Avrilia Floratou (Microsoft)
Jana Giceva (Technical University of Munich)
Katja Hose (Technical University of Vienna)
H. V. Jagadish (University of Michigan)
Panos Kalnis (King Abdullah University of Science and Technology)
Georgia Koutrika (Athena Research Center)
Eric Lo (Chinese University of Hong Kong)
Nikos Mamoulis (University of Ioannina)
Stefan Manegold (CWI)
Ioana Manolescu (Inria and Polytechnic Institute of Paris)
Norman May (SAP SE)
Umar Farooq Minhas (Apple)
Fatemeh Nargesian (University of Rochester)
Beng Chin Ooi (National University of Singapore)
Fatma Ozcan (Google)
Tamer Ozsu (University of Waterloo)
Tilmann Rabl (Hasso Plattner Institute and University of Potsdam)
Mirek Riedewald (Northeastern University)
Jennie Rogers (Northwestern University)
Alkis Simitsis (Athena Research Center)
Letizia Tanca (Polytechnic University of Milan)
Nan Tang (Hong Kong University of Science and Technology (GZ))
Yuanyuan Tian (Microsoft)
Yongxin Tong (Beihang University)
Pinar Tozun (IT University of Copenhagen)
Yannis Velegrakis (Utrecht University)
Matthias Weidlich (Humboldt University of Berlin)
Steven E. Whang (Korea Advanced Institute of Science and Technology)
Raymond Chi-Wing Wong (Hong Kong University of Science and Technology)

Publication Editors
Xiaoou Ding (Harbin Institute of Technology)
Subhadeep Sarkar (Brandeis University)
Giovanni Simonini (University of Modena and Reggio Emilia)

PVLDB Managing Editor
Jun Yang (Duke University)

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

Review Board

Ahmed S. Abdelhamid (Purdue University)
Ziawasch Abedjan (TU Berlin)
Ahmed Aly (Google)
Mohammad Javad Amiri (Stony Brook University)
Yael Amsterdamer (Bar-Ilan University)
Renzo Angles (Universidad de Talca)
Alexander Artikis (University of Piraeus)
Joy Arulraj (Georgia Tech)
Abolfazl Asudeh (University of Illinois Chicago)
Maurizio Atzori (University of Cagliari)
Nikolaus Augsten (University of Salzburg)
Zhifeng Bao (RMIT University)
Ilaria Bartolini (University of Bologna)
Johes Bater (Tufts University)
Lawrence Benson (HPI and University of Potsdam)
Sonia Bergamaschi (University of Modena and Reggio Emilia)
Anna Bernasconi (Politecnico di Milano)
Arnab Bhattacharya (IIT Kanpur)
Alexander Boehm (SAP SE)
Paul Boniol (Universite de Paris)
Renata Borovica-Gajic (University of Melbourne)
Panagiotis Bouros (Johannes Gutenberg University Mainz)
Vanessa Braganholo (Fluminense Federal University)
Matteo Brucato (Microsoft Research)
Michael J. Cahill (University of Sydney)
Diego Calvanese (Free University of Bozen Bolzano)
Jesus Camacho-Rodriguez (Microsoft)
Helena Caminal (Google)
Huiping Cao (New Mexico State University)
Yang Cao (University of Edinburgh)
Zhao Cao (Huawei Technologies)
Zhichao Cao (Arizona State University)
Matteo Ceccarello (University of Padova)
Chengliang Chai (Beijing Institute of Technology)
Yunpeng Chai (Renmin University of China)
Harry Kai-Ho Chan (The University of Sheffield)
Tsz Nam Chan (Shenzhen University)
Subarna Chatterjee (Harvard University)
Cindy Chen (University of Massachusetts Lowell)
Lu Chen (Zhejiang University)
Hong Cheng (The Chinese University of Hong Kong)
Rada Chirkova (NC State University)
Theodoros Chondrogiannis (University of Konstanz)
Shihabur Chowdhury (Apple)
George Christodoulou (TU Delft)
Periklis Chrysogelos (Oracle)
Gao Cong (Nanyang Technological Univesity)
Alex Conway (Cornell Tech)
Andrew Crotty (Northwestern University)
Bin Cui (Peking University)
Patrick Damme (TU Berlin)
Roshan Dathathri (Microsoft Research)
Jesse Davis (MongoDB)
Cagatay Demiralp (MIT)
Dong Deng (Rutgers University New Brunswick)
Laxman Dhulipala (University of Maryland, College Park)
Shimin Di (The Hong Kong University of Science and Technology)
Claudia Diamantini (Universita Politecnica delle Marche)
Anton Dignos (Free University of Bozen Bolzano)
Bailu Ding (Microsoft Research)
Bolin Ding (Alibaba Group)
Jialin Ding (Amazon Web Services)
Anh Dinh (Deakin University)
AnHai Doan (University of Wisconsin Madison)
Christos Doulkeridis (University of Pireaus)
Stefania Dumbrava (ENSIIE)
Ahmed Eldawy (University of California Riverside)
Mohamed Eltabakh (Qatar Foundation)
Venkatesh Emani (Microsoft)
Ju Fan (Renmin University of China)
Zhiwei Fan (Meta)
Yixiang Fang (The Chinese University of Hong Kong)
Anna Fariha (University of Utah)
Ziqiang Feng (Google)
Hakan Ferhatosmanoglu (University of Warwick and Amazon Web Services)
Elena Ferrari (University of Insubria)
Donatella Firmani (Sapienza University)
Peter M. Fischer (University of Augsburg)
George Fletcher (Eindhoven University of Technology)
Juliana Freire (New York University)
Sainyam Galhotra (Cornell University)
Johann Gamper (Free University of Bozen Bolzano)
Yunjun Gao (Zhejiang University)
Paolo Garza (Politecnico di Torino)
Chang Ge (University of Minnesota)
Tingjian Ge (University of Massachusetts Lowell)
Rainer Gemulla (Universitat Mannheim)
Nikos Giatrakos (Technical University of Crete)
Aristides Gionis (KTH Royal Institute of Technology)
Boris Glavic (Illinois Institute of Technology)
Lukasz Golab (University of Waterloo)
Jonathan Goldstein (Microsoft)
Sven Groppe (Universitat zu Lubeck)
Michael Grossniklaus (University of Konstanz)
Anja Gruenheid (Microsoft)
Le Gruenwald (The University of Oklahoma)
Vincenzo Gulisano (Chalmers University of Technology)
Rihan Hai (TU Delft)
Wook-Shin Han (POSTECH)
Mohamed S. Hassan (Oracle)
Oktie Hassanzadeh (IBM Research)
Wenjia He (University of Michigan)
Xi He (University of Waterloo)
Yeye He (Microsoft Research)
Meichun Hsu (Oracle)
Haibo Hu (The Hong Kong Polytechnic University)
Xiao Hu (University of Waterloo)
Qiang Huang (National University of Singapore)
Xin Huang (Hong Kong Baptist University)
Yan Huang (University of North Texas)
Zi Helen Huang (University of Queensland)
Madelon Hulsebos (University of California Berkeley)
Romain Ilbert (Huawei Paris Research Center)
Matteo Interlandi (Microsoft)

Ekaterini Ioanou (Tilburg University)
Gabriela Jacques-Silva (Facebook)
Fuad Jamour (Amazon Web Services)
Soren Kejser Jensen (Aalborg University)
Peiquan Jin (University of Science and Technology of China)
Alekh Jindal (SmartApps)
Hyungsoo Jung (Seoul National University)
Vasiliki Kalavri (Boston University)
Vana Kalogeraki (Athens University of Economics and Business)
Eser Kandogan (Megagon Labs)
Daniel Kang (UIUC)
Zoi Kaoudi (IT University of Copenhagen)
Pinar Karagoz (Middle East Technical University (METU))
Bojan Karlas (Harvard University)
Asterios Katsifodimos (TU Delft)
Oliver A. Kennedy (University at Buffalo SUNY)
Arijit Khan (Aalborg University)
Guy Khazma (University of Toronto)
Haridimos Kondylakis (FORTH-ICS)
Arnd Christian Konig (Microsoft)
Chrysanthi Kosyfaki (The University of Hong Kong)
Nick Koudas (University of Toronto)
Paraschos Koutris (University of Wisconsin Madison)
Mayuresh Kunjir (Amazon Web Services)
Alexandros Labrinidis (University of Pittsburgh)
Wolfgang Lehner (TU Dresden)
Chuan Lei (Amazon Web Services)
Viktor Leis (TU Munich)
Alberto Lerner (University of Fribourg)
Ulf Leser (Humboldt-Universitat zu Berlin)
Guoliang Li (Tsinghua University)
Jia Li (The Hong Kong University of Science and Technology (GZ))
Jianxin Li (Deakin University)
Tian Li (Carnegie Mellon University)
Tianyu Li (MIT)
Yinan Li (Microsoft Research)
Yuchen Li (Singapore Management University)
Xiang Lian (Kent State University)
Shen Liang (Universite Paris Cite)
Michele Linardi (CYU)
Matteo Lissandrini (University of Verona)
Chunwei Liu (MIT)
Jinfei Liu (Zhejiang University)
Xueli Liu (Tianjin University)
Cheng Long (Nanyang Technological University)
Baotong Lu (Microsoft Research)
Jiaheng Lu (University of Helsinki)
Siqiang Luo (Nanyang Technological University)
Yuyu Luo (The Hong Kong University of Science and Technology (GZ))
Manisha Luthra (TU Darmstadt)
Joana M. F. da Trindade (MIT)
Chenhao Ma (The Chinese University of Hong Kong)
Lin Ma (University of Michigan)
Amr Magdy (University of California Riverside)
Ahmed Mahmood (Google)
Sujaya Maiyya (University of Waterloo)
Neha Makhija (Northeastern University)
Silviu Maniu (Universite Grenoble Alpes)
Essam Mansour (Concordia University)
Ryan Marcus (University of Pennsylvania)
Amelie Marian (Rutgers University)
Davide Martinenghi (Politecnico di Milano)
Venkata Vamsikrishna Meduri (IBM Research - Almaden)
Sharad Mehrotra (University of California Irvine)
Alexandra Meliou (University of Massachusetts Amherst)
Paolo Merialdo (Universita degli Studi Roma Tre)
Amine Mhedhbi (Polytechnique Montreal)
Xiaoye Miao (Zhejiang University)
Sebastian Michel (RPTU Kaiserslautern Landau)
Katsiaryna Mirylenka (IBM Research Zurich)
Madhulika Mohanty (Inria Saclay)
Mohamed Mokbel (University of Minnesota Twin Cities)
Mirella M. Moro (Universidade Federal de Minas Gerais)
Davide Mottin (Aarhus University)
Kyriakos Mouratidis (Singapore Management University)
Ingo Müller (Google)
Balakrishnan Narayanaswamy (Amazon)
Mario Nascimento (Northeastern University)
Parimarjan Negi (MIT)
Quoc Viet Hung Nguyen (Griffith University)
Milos Nikolic (University of Edinburgh)
Matthaios Olma (MongoDB)
Prashant Pandey (University of Utah)
George Papadakis (University of Athens)
Dimitris Papadias (The Hong Kong University of Science and Technology)
Odysseas Papapetrou (TU Eindhoven)
John Paparrizos (The Ohio State University)
George Papastefanatos (ATHENA Research Center)
Stefano Paraboschi (Universita degli Studi di Bergamo)
Aditya Parameswaran (University of California Berkeley)
Yongjoo Park (UIUC)
Eliana Pastor (Politecnico di Torino)
Jignesh Patel (Carnegie Mellon University)
Marco Patella (University of Bologna)
Torben Bach Pedersen (Aalborg University)
Botao Peng (Chinese Academy of Sciences)
Peng Peng (Hunan University)
Matthew J. Perron (MIT)
Ilia Petrov (Reutlingen University)
Holger Pirk (Imperial College)
Stefan Plantikow (Neo4j)
Orestis Polychroniou (Amazon)
Danica Porobic (Oracle)
Abdulhakim Qahtan (Utrecht University)
Abdul Quamar (Google)
Weixiong Rao (Tongji University)
Berthold Reinwald (IBM Research Almaden)
El Kindi Rezig (MIT)
Daniel Ritter (SAP)
Oscar Romero (Universitat Politecnica de Catalunya)
Kexin Rong (Georgia Institute of Technology)
Abhishek Roy (Snowflake)

Florin Rusu (University of California Merced)
Sourav S. Bhowmick (Nanyang Technological University)
Ibrahim Sabek (University of Southern California)
Mohammad Sadoghi (University of California Davis)
Semih Salihoglu (University of Waterloo)
Maria Luisa Sapino (University of Torino)
Subhadeep Sarkar (Brandeis University)
Kai-Uwe Sattler (TU Ilmenau)
Patrick Schafer (Humboldt-Universitat zu Berlin)
Felix M. Schuhknecht (Johannes Gutenberg University Mainz)
Maximilian E. Schule (University of Bamberg)
Malte Schwarzkopf (Brown University)
Rathijit Sen (Microsoft)
Jiwon Seo (Seoul National University)
Juan Sequeda (data.world)
Marco Serafini (University of Massachusetts Amherst)
Amir Shaikhha (University of Edinburgh)
Shantanu Sharma (New Jersey Institute of Technology)
Yanyan Shen (Shanghai Jiao Tong University)
Jieming Shi (The Hong Kong Polytechnic University)
Roee Shraga (WPI)
Tarique Siddiqui (Microsoft Research)
Giovanni Simonini (University of Modena and Reggio Emilia)
Utku Sirin (Harvard University)
Spiros Skiadopoulos (University of the Peloponnese)
Dimitrios Skoutas (Athena Research Center)
Shaoxu Song (Tsinghua University)
Divesh Srivastava (AT&T Chief Data Office)
Chrysoula Stathakopoulou (Chainlink Labs)
Kostas Stefanidis (Tampere University)
Kurt Stockinger (ZHAW Zurich University of Applied Sciences)
Uta Storl (University of Hagen)
Shixuan Sun (Shanghai Jiao Tong University)
Ki Hyun Tae (KAIST)
Dixin Tang (University of Texas Austin)
Jing Tang (The Hong Kong University of Science and Technology (GZ))
Mingjie Tang (Sichuan University)
Bo Tang (Southern University of Science and Technology)
Egemen Tanin (University of Melbourne)
Ernest Teniente (Universitat Politecnica de Catalunya)
Arash Termehchy (Oregon State University)
Jens Teubner (TU Dortmund)
Riccardo Torlone (Roma Tre University)
Goce Trajcevski (Iowa State University)
Immanuel Trummer (Cornell University)
Eleni Tzirita Zacharatou (IT University of Copenhagen)
Katerina Tzompanaki (CY Cergy Paris University)
Leong Hou U (University of Macau)
Alexander van Renen (UTN)
Genoveva Vargas-Solar (CNRS LIRIS)
Nalini Venkatasubramanian (University of California Irvine)
Hannes Voigt (Neo4j)
Hongzhi Wang (Harbin Institute of Technology)
Ning Wang (Beijing Jiaotong University)
Qitong Wang (Universite Paris Cite)
Sibo Wang (The Chinese University of Hong Kong)
Tianzheng Wang (Simon Fraser University)
Yifan Wang (University of Florida)
Sai Wu (Zhejiang University)
Yinghui Wu (Case Western Reserve University)
Yuncheng Wu (Renmin University of China)
Xiaokui Xiao (National University of Singapore)
Jianliang Xu (Hong Kong Baptist University)
Jianqiu Xu (Nanjing University of Aeronautics and Astronautics)
Nikolay Yakovets (TU Eindhoven)
Xiao Yan (Centre for Perceptual and Interactive Intelligence (CPII))
Hongzhi Yin (The University of Queensland)
Man Lung Yiu (The Hong Kong Polytechnic University)
Brit Youngmann (Technion)
Jeffrey Xu Yu (The Chinese University of Hong Kong)
Xiaohui Yu (York University)
Yi Yu (NII)
Ye Yuan (Beijing Institute of Technology)
Cong Yue (National University of Singapore)
Demetrios Zeinalipour-Yazti (University of Cyprus)
Yuxiang Zeng (Beihang University)
Steffen Zeuch (TU Berlin)
Chao Zhang (University of Waterloo)
Chen Zhang (The Hong Kong Polytechnic University)
Huanchen Zhang (Tsinghua University)
Meihui Zhang (Beijing Institute of Technology)
Minjia Zhang (Microsoft AI and Research)
Qizhen Zhang (University of Toronto)
Xiaofei Zhang (University of Memphis)
Yanfeng Zhang (Northeastern University)
Bo Zhao (Aalto University)
Zhuoyue Zhao (University at Buffalo)
Bolong Zheng (Huazhong University of Science and Technology)
Kaiping Zheng (National University of Singapore)
Jingren Zhou (Alibaba Group)
Xuan Zhou (East China Normal University)
Yongluan Zhou (University of Copenhagen)
Yiwen Zhu (Microsoft)
Jia Zou (Arizona State University)
Lei Zou (Peking University)
Kostas Zoumpatianos (Snowflake)
Andreas Zufle (Emory University)

NDUSTRIAL TRACK CHAIRS AND REVIEWERS - Vol. 18

Industrial Track Program Chairs
Surajit Chaudhuri (Microsoft, USA)
Nikos Ntarmos (Huawei Europe, UK)
Jingren Zhou (Alibaba, China)

Industrial Track Program Committee Members
Anastasia Ailamaki (EPFL/RAW Labs, Switzerland)
Amirhossein Aleyasen (Datometry, USA)
Gustavo Alonso (ETH (University), Switzerland)
Lyublena Antova (ex-Meta (now: Antova.Rocks, USA)
Angela Bonifati (Lyon 1 University, France)
Peter Boncz (CWI (University), Netherlands)
Nico Bruno (Microsoft, USA)
Jianjun Chen (ByteDance, USA)
Chen Cheng (ByteDance, Singapore)
Sudipto Das (AWS, USA)
Amol Deshpande (Aryn and University of Maryland College Park USA)
Bolin Ding (Alibaba, China)
AnHai Doan (University of Wisconsin (Madison, USA)
Mohamed Eltabakh (Qatar Foundation (University), Qatar)
Johannes Gehrke (Microsoft, USA)
Ihab Ilyas (Apple, USA)
Konstantinos Karanasos (Meta, USA)
Hanuma Kodavalla (Microsoft, USA)
Georgia Koutrika (Athena RC (University), Greece)
Justin Levandoski (Google, USA)
Allison Lee (Snowflake, USA)
Feifei Li (Alibaba, China)
Guoliang Li (Tsinghua University, China)
Norman May (SAP, Germany)
C. Mohan (Hong Kong Baptist University and Tsinghua University, China)
Lev Novik (Databricks, USA)
Beng Chin Ooi (National University of Singapore, Singapore)
Ippokratis Pandis (AWS, USA)
Krishna Kantikiran Pasupuleti (Oracle, USA)
Jignesh Patel (CMU (University), USA)
Stavros Papadopoulos (TileDB, USA)
Danica Porobic (Oracle, Switzerland)
Rajeev Rastogi (Amazon (Retail), India)
Yannis Sismanis (Databricks, USA)
Rebecca Taft (Cockroach Labs, USA)
Efthymia Tsamoura (Samsung, UK)
Stratis Viglas (Google, UK)
Jim Webber (Neo4j, UK)
Quanqing Xu (Oceanbase, China)
Kai Zeng (Huawei, China)

DEMONSTRATION TRACK CHAIRS AND REVIEWERS - Vol. 18

Demonstration Track Program Chairs

Sourav S Bhowmick (Nanyang Technological University, Singapore)
Philippe Bonnet (University of Copenhagen, Denmark)

Demonstration Track Associate Editors

Aditya Parameswaran (University of California, Berkeley)
Bjorn Jonsson (Reykjavik University)
Byron Choi (Hong Kong Baptist University)
Curtis Dyreson (Utah State University)
Danica Porobic (Oracle)
Felix Naumann (Hasso Plattner Institute, University of Potsdam)
George Fletcher (Eindhoven University of Technology the Netherlands)
Jana Giceva (TU Munich)
Melanie Herschel (Nanyang Technological University)
Nan Tang (HKUST, GZ)
Sebastian Link (University of Auckland)
Stefanie Scherzinger (University of Passau)
Steffen Zeuch (TU Berlin)
Tianzheng Wang (Simon Fraser University)

Demonstration Track Program Committee Members

Ali Hadian (Imperial College London)
Andra-Denis Ionescu (TU Delft)
Anna Fariha (University of Utah)
Arash Termehchy (Oregon State University)
Bo Tang (Southern University of Science and Technology)
Boris Glavic (University of Illinois Chicago)
Carlo Sartiani (Università della Basilicata)
Chuan Xiao (Osaka University, Nagoya University)
Dong Deng (Rutgers University)
Fabian Panse (University of Augsburg)
Guozhong Li (King Abdullah University of Science & Technology)
Gylfi Guðmundsson (Reykjavik University)
Hazar Harmouch (University of Amsterdam)
Hui Li (Xidian University)
Ibrahim Sabek (University of Southern California)
Irena Holubova (Charles University)
Jan Hidders (Birkbeck, University of London)
Jiaxin Jiang (National University of Singapore)
Jiwon Seo (Seoul National University)
John Paparrizos (The Ohio State University)
Kostas Stefanidis (Tampere University)
Kyoungmin Kim (EPFL)
Larissa Capobianco Shimomura (University of Stuttgart)
Leong Hou U (University of Macau)
Lorena Etcheverry (Instituto de Computación, Facultad de Ingeniería, Universidad de la República)
Lucas Braun (Oracle Labs)
Madelon Hulsebos (CWI)
Madhulika Mohanty (Inria Saclay)
Manisha Luthra (TU Darmstadt)
Marios Fragkoulis (TU Delft)
Martin Hentschel (IT University of Copenhagen)
Michael Mior (Rochester Institute of Technology)
Nikos Bikakis (Athena)
Odysseas Papapetrou (TU Eindhoven)
Oscar Romero (Universitat Politècnica de Catalunya)
Petra Selmer (Bloomberg)
Raja Appuswamy (Eurecom)
Renata Borovica-Gajic (University of Melbourne)
Rihan Hai (TU Delft)
Sajjadur Rahman (Megagon Labs)
Shaleen Deep (Microsoft Gray Systems Lab)
Stefania Dumbrava (ENSIIE)
Supun Nakandala (Databricks)
Tarique Siddiqui (Microsoft Research)
Tobias Ziegler (Technische Universität München)
Toshiyuki AMAGASA (University of Tsukuba)
Uta Störl (University of Hagen)
Verena Kantere (University of Ottawa)
Viktor Sanca (EPFL)
Vivek Shah (Samsung)
Yongluan Zhou (University of Copenhagen)
Zhengjie Miao (Simon Fraser University)

TUTORIAL TRACK CHAIRS AND REVIEWERS - Vol. 18

Tutorial Track Program Chairs
Hakan Ferhatosmanoglu (University of Warwick and AWS)
Madelon Hulsebos (CWI)

Tutorial Track Program Committee Members
Aditya Parameswaran (University of California, Berkeley)
Andreas Kipf (UTN)
Cheng Long (Nanyang Technological University)
Fatma Ozcan (Google)
Gerardo Vitagliano (MIT CSAIL)
Jianguo Wang (Purdue University)
Matteo Interlandi (Microsoft)
Matthias Boehm (TU Berlin)
Sharad Mehrotra (UC Irvine)
Utku Sirin (Harvard University)
Wenjie Zhang (University of New South Wales)
Xiaofang Zhou (HKUST)
Yang Cao (University of Edinburgh)
Yunyao Li (Adobe)
Zoi Kaoudi (IT University of Copenhagen)

LETTER FROM THE EDITORS IN CHIEF

It is our pleasure to present the twelfth (August 2025) issue of PVLDB Volume 18. While the first eleven issues covered research track papers accepted to PVLDB and presented at the VLDB 2025 Conference in London, Issue 12 includes all other contributions, which were also part of the VLDB 2025 program: the VLDB Endowment Award talks, the Keynote talks, Panels, Workshops, as well as peer-reviewed Industrial papers, Demonstrations, and Tutorials.

The Industrial Track of VLDB 2025 invited papers on industrial products and services as well as industrial prototypes. Our call for papers discouraged submission of academic projects that are well served by the research and the demonstration tracks of the conference. Unlike past years, we had only a "light" revision round intended primarily to improve presentation. We accepted 34 papers from among the 107 papers that were submitted. Each paper was reviewed by at least two members of the industrial track program committee and some of the papers were assigned additional reviewers as needed. The papers have been organized in six sessions: Distributed Systems, Data Platforms for Analytics, Document/Graph/Vector Databases, Machine Learning/AI & Databases, Transactions & Concurrency Control, and Database Engines.

The Demonstration Track of VLDB 2025 continues to serve as a prominent platform for disseminating and exhibiting the latest innovations in the field of data management. To improve the quality and fairness of the review process at scale, a group of distinguished and diverse researchers was appointed as Associate Editors (AEs), who were entrusted with overseeing review quality and facilitating reviewer discussions. In addition, automated conflict-of-interest (CoI) checks were conducted using the CLOSET System. Furthermore, the review form was specifically tailored to better align with the unique characteristics of demonstration submissions. The track received a substantial number of submissions, totaling 150, 55 of which were accepted for presentation, resulting in an acceptance rate of 36.7%. Each submission was reviewed by a minimum of two expert reviewers under the responsibility of an AE. The candidates for the Best Demonstration Award were proposed by AEs on the basis of the accepted proposals.

The Tutorial Track of VLDB 2025 covers state-of-the-art research, development, and applications in data management or related areas, including interdisciplinary areas. This year, we received 25 submissions, out of which we accepted 13 tutorials. We thank all authors for their contributions to the track. The Tutorial Program Committee was assembled with the diversity of topics and demographics within our community in mind. Each paper was assigned to the committee and subsequently reviewed by the Tutorial Program Chairs for a final decision, in coordination with the Program Committee Chairs and the General Chairs. The accepted proposals cover a wide range of exciting topics, including learned query processing, LLM inference and model management, natural language interfaces to databases, data discovery in data lakes, synthetic data, vector search, graph databases and analytics.

VLDB 2025 has a rich panel program with the following 4 panels:

- Neural Relational Data: Tabular Foundation Models, LLMs... or both? (Moderators: Paolo Papotti, Carsten Binnig; Panelists: Floris Geerts, Johannes Hoffart, Madelon Hulsebos, Fatma Özcan, Gael Varoquaux)
- Beyond Incrementalism: How to Change the World Through Data Systems Research (Moderator: Viktor Leis; Panelists: Anastasia Ailamaiki, Peter Boncz, Badrish Chandramouli, Andy Pavlo)
- Where Does Academic Database Research Go From Here? (Moderators: Eugene Wu, Raul Castro Fernandez; Panelists: Shreya Shankar, Natacha Crooks, Jiannan Wang, Gustavo Alonso, Divesh Srivastava)
- Open Science: A New Paradigm for the Research Lifecycle (Moderator: Yannis Ioannidis; Panelists: Yanlei Diao, Dame Wendy Hall, Wolfgang Lehner, Natalia Manola, Julia Stoyanovich)

VLDB 2025 received 24 workshop proposals. After reviewing the proposals and evaluating their merits based on their history, relevance of topics, and distinctiveness, we selected 14 high-quality workshops (8 x full-day and 6 x half-day). Like previous years, we erred on the side of accepting a larger number of workshops to ensure a good diversity of topics and allow new and revived workshops to establish themselves with sufficiently many submissions. This year's conference program included the following workshops:

Monday, September 1:

- ADMS: 16th International Workshop on Accelerating Analytics and Data Management Systems Using Modern Processor and Storage Architectures (Chairs: Rajesh Bordawekar, Tirthankar Lahiri)
- AIDB: 6th International Workshop on Applied AI for Database Systems and Applications (Chairs: Thaleia Doudali, Subru Krishnan, Umar Farooq Minhas)
- LLM+Graph: 2nd International Workshop on Data Management Opportunities in Bringing LLMs with Graph Data (Chairs: Yixiang Fang, Arijit Khan, Tianxing Wu, Da Yan)
- QDB: 14th International Workshop on Quality in Databases (Chairs: Lisa Ehrlinger, Lorena Etcheverry, Hazar Harmouch)
- TPCTC: 17th TPC Technology Conference on Performance Evaluation and Benchmarking (Chairs: Raghunath Nambiar, Meikel Poess)

Friday, September 5:

- CDMS: 3rd International Workshop on Composable Data Management Systems (Chairs: Satyanarayana R. Valluri, Mohamed Zait)
- DaSH: 6th Workshop on Data Science with Human in the Loop (Chairs: Eduard Dragut, Yunyao Li, Lucian Popa, Kun Qian, Sherry Tongshuang Wu)
- DATAI: 2nd International Workshop on Data-driven AI (Chairs: Nan Tang, Hongzhi Wang, Lei Cao, Chengliang Chai, Xiaoou Ding)
- DEC: 3rd Data EConomy Workshop (Chairs: Santiago Andrés, George Konstantinidis)
- Guide-AI: 2nd Workshop on Governance, Understanding, and Integration of Data for Effective and Responsible AI (Chairs: Babak Salimi, Sainyam Galhotra)
- LLM+Spatial: 1st Workshop on Large Language Models for Spatial-rich Data Management (Chairs: Cheng Long, Bernhard Seeger, Yongxin Tong, Jianqiu Xu)
- LSGDA: 4th International Workshop on Large-Scale Graph Data Analytics (Chairs: Wenjie Zhang, Ying Zhang, Wentao Li, Dong Wen, Zhengyi Yang)
- LS-NSL: 1st Workshop on New Ideas for Large-Scale Neurosymbolic Learning Systems (Chairs: Efthymia Tsamoura, Pablo Barceló, Jacopo Urbani)
- TaDA: 3rd International Workshop on Tabular Data Analysis (Chairs: Vasilis Efthymiou, Oktie Hassanzadeh, Sainyam Galhotra, Ernesto Jiménez-Ruiz)

The 2025 edition of VLDB's PhD Workshop features 15 student presentations, a keynote by Prof. Gerhard Weikum on "Adventure and Beauty in Data Systems Research", and a panel discussion.

VLDB 2025 features 3 Keynote talks:

- Alphabets, Grammars, Calculators, and the End of Hand-Crafted Systems (Prof. Stratos Idreos, Harvard University)
- Bridging Disciplines in Data Management Research to Solve Complex Data Problems (Prof. Juliana Freire, New York University)
- Bringing the Operational and Analytical Worlds Together with Lakebase (Prof. Matei Zaharia, University of California Berkeley and Databricks)

Besides the Best Paper and Best Demonstration awards announced during the conference, VLDB 2025 features 3 VLDB Endowment awards with invited talks in the program:

- 2025 VLDB Test of Time Award: Viktor Leis (TU Munich), Andrey Gubichev (TU Munich), Atanas Mirchev (TU Munich), Peter Boncz (CWI), Alfons Kemper (TU Munich), Thomas Neumann (TU Munich) for their VLDB 2015 paper "How Good Are Query Optimizers, Really?"
- 2025 VLDB Women in Database Research Award: Angela Bonifati (Lyon 1 University) for "pioneering contributions in graph databases, and leadership in the international database community"
- 2025 VLDB Early Career Research Contribution Award: Xiangyao Yu (University of Wisconsin-Madison) for "significant contributions in scalable transaction management"

In closing, we wish to collectively express our deep gratitude to all members of various program committees as well as our proceedings chairs who worked tirelessly in the past year to ensure the timely and smooth publication of PVLDB Volume 18.

Themis Palpanas and Nesime Tatbul

Editors-in-Chief of PVLDB Vol. 18
Program Chairs for VLDB 2025

Surajit Chaudhuri, Nikos Ntarmos, and Jingren Zhou
Industrial Program Chairs

Sourav S. Bhowmick and Philippe Bonnet
Demonstration Program Chairs

Hakan Ferhatosmanoglu and Madelon Hulsebos
Tutorial Program Chairs

Jana Giceva and Alexandra Meliou
Panel Program Chairs

John Paparrizos and Norman Paton
Workshop Program Chairs

Sonia Bergamaschi and Raul Castro Fernandez
PhD Workshop Chairs

Jiuqi Wei
Workshop Proceedings Chair

Xiaoou Ding, Subhadeep Sarkar, and Giovanni Simonini
Proceedings Chairs